



(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
27.06.2001 Bulletin 2001/26

(51) Int Cl.7: **H04N 1/40**

(21) Application number: 00311430.3

(22) Date of filing: 20.12.2000

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
 MC NL PT SE TR**
 Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 20.12.1999 US 172780 P
29.06.2000 US 215000 P
29.06.2000 US 214951 P
04.08.2000 US 632543 P

(71) Applicant: **Texas Instruments Incorporated**
Dallas, Texas 75251 (US)

(72) Inventors:

- Hung, Ching-Yu
Piano, Texas 75024 (US)
- Tamama, Hideo
Tsukuba, Ibaraki (JP)
- Osamoto, Akira
Inashiki, Ibaraki (JP)
- Koshiba, Osamu
Tsukuba, Ibaraki (JP)

- Yamauchi, Satoru
Tsuchiura, Ibaraki (JP)
- Zhou, Minhua
Plano, Texas 75093 (US)
- Ilgner, Klaus
81549 München (DE)
- Talluri, Rajendra
Plano, 75203 (US)
- Yoo, Youngjun
Richardson, Texas 75080 (US)
- Liang, Jie
Plano, Texas 75025 (US)
- Tsai, Mandy
Taipei 111 (KR)
- Tsunoda, Kiyoshi
Tokohama-city, Kanagawa (JP)
- Inamori, Shinri
Yokohama-shi, Kanagawa (JP)

(74) Representative: **Potter, Julian Mark et al**
D. Young & Co.,
21 New Fetter Lane
London EC4A 1DA (GB)

(54) **Digital camera system and method**

(57) Digital Still Camera (DSC) includes separate preview engine, burst mode compression/decompression engine, image pipeline, CCD plus CCD controller.

and memory plus memory controller. ARM microprocessor and DSP share control. Color filter array interpolation by use of green high-frequency added to red and blue plus interpolation.




FIG. 1a

Description

[0001] This application claims priority from provisional applications Serial Nos 60/172.780, filed 20 December 1999; 60/176.272, filed 14 January 2000; 60/177.432, filed 21 January 2000; 60/214.951, filed 29 June 2000 and 60/215.000, filed 29 June 2000.

[0002] This invention relates to integrated circuits, and more particularly, but not exclusively, to integrated circuits and methods for use with digital cameras.

[0003] Recently, Digital Still Cameras (DSCs) have become a very popular consumer appliance appealing to a wide variety of users ranging from photo hobbyists, web developers, real estate agents, insurance adjusters, photo-journalists to everyday photography enthusiasts. Recent advances in large resolution CCD arrays coupled with the availability of low-power digital signal processors (DSPs) has led to the development of DSCs that come quite close to the resolution and quality offered by traditional film cameras. These DSCs offer several additional advantages compared to traditional film cameras in terms of data storage, manipulation, and transmission. The digital representation of captured images enables the user to easily incorporate the images into any type of electronic media and transmit them over any type of network. The ability to instantly view and selectively store captured images provides the flexibility to minimize film waste and instantly determine if the image needs to be captured again. With its digital representation the image can be corrected, altered, or modified after its capture. See for example, Venkataraman et al, "Next Generation Digital Camera Integration and Software Development Issues" in Digital Solid State Cameras: Design and Applications, 3302 Proc. SPIE (1998). Similarly, USP 5,528,293 and USP 5,412,425 disclose aspects of digital still camera systems including storage of images on memory cards and power conservation for battery-powered cameras.

[0004] Viewed from a first aspect, the invention provides a digital still camera architecture with features selected from programmable camera functions plus short audio-video clip capabilities, dual processors plus an image coprocessor, burst mode compression/decompression engine, programmable preview engine, and integration of all camera peripherals including IrDA, USB, NTSC/PAL encoder, DACs for RGB, UART, and compact flash card/smart media card interface.

[0005] This has advantages including flexibility, adaptability, and efficiency.

[0006] Embodiments in accordance with the invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Figures 1a-1b show a preferred embodiment system in functional block format.

Figures 2-6 illustrate data flows.

Figures 7a-7b show CFA arrangements.

Figure 8 is a functional diagram for white balance.

Figures 9a-9b show gamma correction.

Figure 10 illustrates CFA interpolation.

Figures 11a-11b show color conversion.

Figures 12a-12b show a memory controller data flow.

Figure 13 is a functional block diagram of a burst compression/-decompression engine.

Figure 14 is a functional block diagram of a preview engine.

Figure 15 is an on screen display block diagram.

Figure 16 is an on screen display window.

Figure 17 shows a hardware cursor.

Figure 18 illustrates a DSP subsystem.

Figure 19 shows parallel multiply-accumulate datapath.

Figure 20 shows a coprocessor architecture.

Figure 21 illustrates a look-up table accelerator.

Figure 22 is a block diagram of a variable length coder.

Figure 23 shows a bridge.

Figure 24 shows multiprocessor debugging support.

Figure 25 illustrates UART connections.

Figure 26 is a block diagram of flash card/smart card interface.

Figures 27-38 illustrate color filter array interpolations.

Figures 39a-39b and 40 show white balancing.

Figures 41a-41b and 42a-42e indicate image resizing.

Figures 43-45 illustrate tone-scaling.

Figures 46a-46b and 47-48 show synchronization.

Figures 49-52 show decoding buffering.

System overview

[0007] Figures 1a-1b show the various high-level functional blocks in a preferred embodiment digital still camera (DSC) and systems with Figure 1b providing more detail than Figure 1a. In particular, preferred embodiment integrated circuit 100 includes the following items: CCD Controller 102 interfaced with either CCD or CMOS imager 150; preview engine block 104 to convert the data from CCD controller 102 into a format suitable for display using NTSC encoder 106 or a digital LCD interface; burst mode compression-decompression engine 108 to compress the raw image data from CCD controller 102 using a lossless (or lossy, as selected by the user) compression and to write the compressed data to external SDRAM 160 via SDRAM controller 110. This data can then be decompressed by the decompression engine under DSP 122 control, processed, and displayed or stored back to SDRAM 160. DSP subsystem block 120 (DSP 122 and iMX 124 plus Variable Length Coder 126 and buffers 128) performs all the processing of the image data in the capture mode. The data is fetched from SDRAM 160 into image buffer 128 by DSP 122 through requests to SDRAM controller 110, and DSP 122 performs all the image processing and compression required in the capture mode. The Image Extension processor (iMX) 124 acts as a dedicated accelerator to DSP 122 to increase the performance of DSP 122 for the imaging applications.

[0008] RISC microprocessor subsystem (ARM 130 plus memory 132) supports the in-camera Operating Systems (OS). Various OSes and other real-time kernels such as VxWorks, MicroItron, Nucleus, and PSOS may be supported on circuit 100.

[0009] SDRAM controller block 110 acts as the main interface between SDRAM 160 and all the function blocks such as the processors (ARM 130, DSP 122), CCD controller 102, TV encoder 106, preview engine 104, etc. SDRAM controller 110 may support up to 80 MHz SDRAM timing and also provide a low overhead for continuous data accesses. It also has the ability to prioritize the access units to support the real-time data stream of CCD data in and TV display data out.

[0010] Camera shot-to-shot delay is the time it takes for DSC engine 100 to read the data from CCD 150, process it and write it to SDRAM 160. The processing includes the image pipeline stages and also JPEG compression.

[0011] In order to support real-time preview, DSC engine 100 will set CCD 150 in "fast readout" mode, process the data, convert the data to NTSC format, and display the data on a built-in LCD screen (not shown in Figure 1) or TV monitor as the case may be.

[0012] Auto focus, auto exposure and auto white balance (the 3A functions) are performed by DSP 122 while DSC 100 is in the preview mode of operation. DSP 122 reads the image data from SDRAM 160, performs the 3A functions in real-time. The algorithms for the 3A functions are programmable.

[0013] Both interface and progressive CCD and CMOS imagers 150 interface directly to DSC engine 100 using the built-in CCD/CMOS controller 102.

[0014] In-camera operating systems such as MicroItron will be supported efficiently on ARM processor 130 in DSC engine 100. DSC engine 100 also has the capability to support capturing of a rapid sequence of images in the "burst mode" of operation. Bursts at up to 10 frames/sec of 2 Megapixel images will be supported. The duration of the burst sequence is only limited by the size of SDRAM 160 of the DSC system. Also, MPEG compression may be used for short clips. And capabilities for playback of audio-video include circular buffering.

[0015] DSC circuit 100 also includes I/O block 140 with USB core 142 for programming and interrupt processing with ARM 130.

[0016] CCD module 150 includes a CCD imager to sense the images, driver electronics and a timing generator for the necessary signals to clock the CCD, correlated double sampling and automatic gain control electronics. This CCD data is then digitized and fed into the DSC Engine 100.

[0017] SDRAM 160 may be any convenient size and speed SDRAM.

[0018] DSC systems may be even more versatile with the ability to annotate images with text/speech. The preferred embodiment programmable DSP allows easy inclusion of a modem and/or a TCP/IP Interface for direct connection to the Internet. DSCs may run complex multi-tasking operating systems to schedule the various real-time tasks.

[0019] Thus the preferred embodiments provide platforms for programmable camera functions, dual processors (ARM and DSP) plus an image coprocessor, burst mode compression/decompression engine, programmable preview engine, and integration of all camera peripherals including IrDA, USB, NTSC/PAL encoder, DACs for RGB, UART, and compact flash card/smart media card interface. Further, the platforms can provide both camera functions and digital audio playback on the same integrated circuit.

[0020] The following sections provide more detail of the functions and modules.

DSC operating modes

[0021] The preferred embodiment systems have (1) Preview mode, (2) Capture mode, (3) Playback mode, and (4) Burst mode of operation as follows.

(1) Preview mode has data flow as illustrated in Figure 2. ARM 130 sets CCD 150 into high-frame-rate readout mode (reduced vertical resolution). ARM 130 enables preview engine 104 and sets the appropriate registers for the default parameters. The raw CCD data is streamed into preview engine 104 and, after preview engine processing, is streamed into SDRAM 160. ARM 130 enables TV encoder 106 to display the preview engine output. Preview engine 104 processing (hardware) includes gain control, white balance, CFA interpolation, down-sampling, gamma correction, and RGB to YUV conversion. ARM 130 commands DSP 122 to perform auto exposure and auto white balance whenever required. DSP 122 processing includes auto exposure, auto white balance, and auto focus. ARM 130 receives new parameters for preview engine 104 and loads the preview engine hardware with these parameters. The output is full resolution CCIR 601 NTSC/PAL and real-time updating of gain, white balance, and auto focus.

(2) Capture mode has data flow as illustrated in Figure 3a. ARM 130 sets CCD 150 in "fine" readout mode, full resolution. The CCD data is read directly into SDRAM 160 through SDRAM controller 110. ARM 130 commands DSP 122 (plus IMX 124 and VLC engine 126) perform capture processing: black clamp, fault pixel correction, shading compensation, white balancing, gamma correction, CFA interpolation, color space conversion, edge enhancement, false color suppression, 4:2:0 down-sampling, and JPEG compression. The DSP stores compressed data in the SDRAM. ARM 130 writes the compressed data to compact flash/smart media 182.

[0022] The computation is scheduled as two threads: iMX on one thread, the other units on the other thread. Figure 3b shows timing and data flow with threads related to buffers A and B.

(3) Playback mode has data flow as illustrated in Figure 4. ARM 130 reads the compressed data from CFC/Smart-media 182 into SDRAM 160 through the SDRAM controller 110 using DMA 162. ARM commands DSP 122 to do "playback". DSP processing (DSP 122 plus IMX 124 and VLC engine 126) includes JPEG decode (bitstream parsing, IDCT, VLD, and down-sampling for aspect ratio) and store uncompressed image data in SDRAM. ARM enables TV encoder 106 to display the image on TV/LCD display. Note that also audio plus video (e.g., MPEG compressed) clips may be played back.

(4) Burst capture mode has data flow as illustrated in Figure 5, and Figure 6 shows offline data processing. ARM 130 sets CCD 150 into fine resolution mode. ARM sets up the burst compression parameters, burst length, number of frames/second, compression ratio (lossy, lossless), etc. ARM enables burst compression engine 108 to write the raw CCD data to SDRAM 160. ARM signals DSP to process each of the stored raw CCD images in the burst. Burst mode decompression engine 108 decompresses each of the burst captured images. DSP processes each of the images as in normal capture and writes the JPEG bitstream to SDRAM 160.

[0023] Burst capture mode is achieved by repeated calls to the regular playback routine with a different JPEG bit-stream each time by ARM 130.

[0024] The preferred embodiment also has MPEG1 capture mode and playback mode.

Image acquisition

[0025] A DSC usually has to perform multiple processing steps before a high quality image can be stored. The first step is the image acquisition. The intensity distribution reflected from the scene is mapped by an optical system onto the imager. The preferred embodiments use CCDs, but a shift to CMOS does not alter the image processing principles. To provide a color image the imager (CCD or CMOS) has each pixel masked by a color filter (such as a deposited dye on each CCD photosite). This raw imager data is normally referred as a Color- Filtered Array (CFA). The masking pattern of the array of pixels in the CCD as well as the filter color primaries vary between different manufactures. In DSC applications, the CFA pattern that is most commonly used is an RGB Bayer pattern that consists of 2x2 cell elements which are tiled across the entire CCD-array. Figure 7a depicts a subset of this Bayer pattern in the matrix block following the CCD camera. Note that half of the pixels are sensitive to green and that the red and blue are balanced to green. Figure 7b shows a subset of the alternative complementary color CFA pattern with yellow, cyan, green, and magenta pixels.

Image pipeline

[0026] CFA data needs to undergo a significant amount of image processing before the image can be finally presented in a usable format for compression or display. All these processing stages are collectively called the "image pipeline". The preferred embodiment DSC may perform multiple processing steps before a high quality image can be stored. Most of the image pipeline processing tasks are multiply-accumulate (MAC) intensive operations, making a DSP a preferred platform. The various image pipeline processing stages are described in the following sections.

A/D converters

[0027] The A/D converter digitizing the CCD imager data may have a resolution of 10 to 12 bits. This allows for a good dynamic range in representing the input image values. Of course, higher resolution implies higher quality images but more computations and slower processing; and lower resolution implies the converse. The A/D converter may be part of the CCD module.

Black clamp

[0028] After A/D conversion the "black" pixels do not necessarily have a 0 value due to a CCD which may still record some current (charge accumulation) at these pixel locations. In order to optimize the dynamic range of the pixel values represented by the CCD imager, the pixels representing black should have a 0 value. The black clamp function adjusts for this by subtracting an offset from each pixel value. Note that there is only one color channel per pixel at this stage of the processing.

Fault pixel interpolation

[0029] CCD-arrays may have defective (missing) pixels, especially arrays with more than 500,000 elements. The missing pixel values are filled by simple interpolation. A high order interpolation may not be necessary because an interpolation is also performed in the CFA interpolation stage. Therefore, the main reason for this preliminary interpolation step is to make the image processing regular by eliminating missing data.

[0030] Typically, the locations of the missing pixels are obtained from the CCD manufacturer. The faulty pixel locations can also be computed by the DSC engine offline. For example, during camera initialization operation, an image with the lens cap closed is captured. The faulty pixels appear as "white spots" while the rest of the image is dark. The faulty pixel locations can then be identified with a simple threshold detector and stored in memory as a bitmap.

[0031] During the normal operation of the DSC the image values at the faulty pixel locations are filled by a simple bilinear interpolation technique.

Lens distortion compensation

[0032] Due to non-linearities introduced by imperfections in lenses, the brightness of the image decreases from the center of the image to the borders of the image. The effects of these lens distortions are compensated by adjustment of the brightness of each pixel as a function of its spatial location. The parameters describing the lens distortions need to be measured with the final system, supported by information supplied by the lens manufacturer.

[0033] The lens adjustment can be accomplished by multiplying the pixel intensity with a constant, where the value of the constant varies with the pixel location. The adjustment needs to be done for both horizontal and vertical directions.

White balance

[0034] White balancing tries to transform the tristimulus values sensed under a certain light condition such that if displayed white appears again as white. In general the colors as captured by the camera do not appear on an output device as they were seen when capturing the scene. A couple of reasons account for that.

[0035] First, the sensitivity of the color filters over the spectral range are slightly different. If exposed with a perfect white light source (constant light spectrum) the tristimulus values sensed by the CCD are slightly different.

[0036] Second, the design of the entire CCD module and the optical system add to the imbalance of the tristimulus values.

[0037] Third, typical illuminants present while recording a scene are not constant. The illuminants have a certain "color", which is typically characterised as "color temperature" (or correlated color temperature). If an image captured under illuminant 1 is displayed under a different illuminant the color appearance changes. This causes a white area to turn a little bit red or a little bit blue.

[0038] Several different approaches for white balancing are known. Most of them multiply the red and blue channels with a factor such that the resulting tristimulus value for a white patch has identical values:

$$\begin{bmatrix} R' \end{bmatrix} = \begin{bmatrix} a1 & 0 & 0 \end{bmatrix} \begin{bmatrix} R \end{bmatrix}$$

$$\begin{bmatrix} G' \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} G \end{bmatrix}, \quad R' = G' = B' \text{ for a neutral (gray) patch}$$

$$\begin{bmatrix} B' \end{bmatrix} = \begin{bmatrix} 0 & 0 & a_2 \end{bmatrix} \begin{bmatrix} B \end{bmatrix}$$

However, as explained later, this approach does not provide correction for changes of the illuminant. Therefore, the white balancing implementation in preferred embodiment system corrects imbalances of the sensor module. The illumination correction is handled at a later stage in the color correction section.

[0039] Typical techniques to calculate the gain factors are

(1) equal energy

$$a_1 = \sum_{(x,y)} g^2(x,y) / \sum_{(x,y)} r^2(x,y)$$

(2) gray world assumption

$$a_1 = \sum_{(x,y)} g(x,y) / \sum_{(x,y)} r(x,y)$$

(3) maximum value in an image is white

$$a_1 = \max_{(x,y)} g(x,y) / \max_{(x,y)} r(x,y)$$

[0040] All of them do not hold in every case. Therefore, by defining the white balancing mainly as a correction of imager module characteristics, the algorithms to obtain the correction values can be made almost scene independent.

[0041] The Figure 8 depicts the simplified realization of the preview engine, giving good results as long as the CCD sensor operates in the linear range. The white balance section below discusses a more sophisticated method.

Gamma correction

[0042] Display devices (TV monitors) used to display images and printers used to print images have a non-linear mapping between the image gray value and the actual displayed pixel intensities. Hence, in the preferred embodiment DSC Gamma correction stage compensates the CCD images to adjust them for eventual display/printing.

[0043] Gamma correction is a non-linear operation. The preferred embodiments implement the corrections as table look ups. The advantages of table look up are high speed and high flexibility. The look-up table data might even be provided by the camera manufacturer.

[0044] With 12-bit data, a full look-up table would have 4K entries, with each entry 8 to 12 bits. For a smaller look-up table, a piecewise linear approximation to the correction curves could be used. For example, the 6 most significant bits could address a 64-entry look-up table whose entries are pairs of values: a base value (8 to 12 bits) and a slope (6 bits). Then the product of the 6 least significant bits and the slope is added to the base value to yield the final corrected value of 8 to 12 bits. Figure 9b illustrates a piecewise linear approximation curve, and Figure 9c the corresponding operations.

[0045] Note that LCD displays can be considered to be linear, making gamma compensation unnecessary. However, LCD display modules usually expect an NTSC input (which is already gamma compensated) and hence perform some "gamma uncorrection" (inverse gamma correction) to compensate for this expected gamma correction. So in the preferred embodiment DSCs using such LCD preview modules, still perform Gamma correction and then NTSC encode the signal before feeding it to the LCD module.

[0046] Gamma correction may be performed at the end of the all the stages of the image pipeline processing and just before going to the display. Alternatively, the image pipeline could perform the Gamma correction earlier in the pipeline: before the CFA interpolation stage.

CFA interpolation

[0047] Due to the use of a color-filtered array (CFA), the effective resolution of each of the color planes is reduced.

At any given pixel location there is only one color pixel information (either of R, G, or B in the case of RGB color primaries). However, it is required to generate a full color resolution (R, G, and B) at each pixel in the DSC. To be able to do this, the missing pixel values (R and B at the G location, etc.) are reconstructed by interpolation from the values in a local neighborhood in the CFA interpolation. To take advantage of the DSP in this system a FIR-kernel is employed as interpolation filter. The length of the filter and the weights vary from one implementation to the other. Also the interband relationship has to be considered. Figure 10 describes the realization of the CFA interpolation in the hardwired preview engine module. It basically employs a 1D FIR kernel for horizontal followed by vertical interpolation.

[0048] The implementation in the DSP subsystem for high quality image processing is different in that it is fully programmable and able to utilize 2D filter kernels. Some background information and a proposal for an improved CFA interpolation technique is given in subsequent sections.

Color correction

[0049] Changes in the color appearance caused by differing illuminants between capture and playback/print cannot be corrected just by balancing the red, green and blue channels independently. To compensate for this, a tone (color) correction matrix maps the RGB pixel values to corrected RGB pixel values that take the illuminant into account.

[0050] The principle is as follows. Let I^1 denote an $N \times N$ diagonal matrix describing the recording illuminant, S the $N \times 3$ matrix denoting the spectral characteristics of the imager module with one column vector for each color, and R the $1 \times N$ column vector describing the reflectance of the scene. The measured tristimulus value X_1 at a pixel location is given by:

$$X_1^T = R^T \cdot I^1 \cdot S$$

Denoting

[0051]

$$SS = S \cdot S^T$$

we can transform the measured tristimulus value X_1 into X_2 , we would have been measured if the scene would have been illuminated by I^2 :

$$X_2^T = X_1^T \cdot S^T \cdot SS^{-1} \cdot I^1 \cdot I^2 \cdot S$$

The 3×3 transform matrix $S^T \cdot SS^{-1} \cdot I^1 \cdot I^2 \cdot S$ can be calculated offline, assuming that the spectral response of the sensor can be measured. Thus it is sufficient to store a set of color correction matrices for different illuminants in the camera.

[0052] Since the subjective preferences of the color appearance changes among users, it is easily possible to include these into the color correction matrix or add a separate step to the image processing pipeline (e.g. "tone scale").

Color space conversion

[0053] After the CFA interpolation and color correction, the pixels are typically in the RGB color space. Since the compression algorithm (JPEG) is based on the YCbCr color space, a color space transformation must be carried out. Also the preferred embodiment DSC generates a NTSC signal output for display on the TV and also to feed into the LCD preview. Hence an RGB to YCbCr color space conversion needs to be carried out. This is a linear transformation and each Y, Cb, Cr value is a weighted sum of the R, G, B values at that pixel location. Figure 11a illustrates the color conversion as realized in the hardwired preview engine. The DSP (playback) implementation is similar in principle but allows a higher precision conversion:

$$\begin{bmatrix} Y \\ Cb \\ Cr \end{bmatrix} = \begin{bmatrix} \alpha_1 & \alpha_2 & \alpha_3 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

$$\begin{bmatrix} Cb \\ Cr \end{bmatrix} = \begin{bmatrix} \alpha_4 & \alpha_5 & \alpha_6 \\ \alpha_7 & \alpha_8 & \alpha_9 \end{bmatrix} \begin{bmatrix} G \\ B \end{bmatrix}$$

Edge enhancement

[0054] After CFA interpolation the images appear a little "smooth" due to the low pass filtering effect of the interpolation filters. To sharpen the images it is sufficient to operate on the Y-component only. At each pixel location we compute the edge magnitude using an edge detector, which is typically a two-dimensional FIR filter. The preferred embodiment uses a 3x3 Laplace-Operator. The edge magnitude is thresholded and scaled and before being added to the original luminance (Y) image to enhance the sharpness of the image.

[0055] The edge enhancement is a high pass filter; this high pass filter also amplifies the noise. To avoid this amplified noise, a threshold mechanism is used to only enhance those portion of the image lying on an edge. The amplitude of the amplified edge may vary. The threshold operation is necessary to reduce amplification of noise. Therefore, only those pixels get enhanced which are an element of an edge. The enhancement signal added to the luminance channel can be represented graphically as in Figure 11b; the parameters t1, t2, and the slope sl can be chosen as seen necessary to obtain the best quality.

False color suppression

[0056] Note that the edge enhancement is only performed in the Y image. At edges the interpolated images of the color channels may not be aligned well. This causes annoying rainbow-like artifacts at sharp edges. Therefore, by suppressing the color components Cb and Cr at edges in the Y-component, these artifacts can be reduced. Depending on the output of the edge detector, the color components Cb and Cr are multiplied by a factor less than 1 on a per pixel basis to suppress the false color artifacts.

Image compression

[0057] The image compression step compresses the image, typically by about 10:1 to 15:1. The preferred embodiment DSC uses JPEG compression. This is a DCT-based image compression technique that gives good performance.

Auto Exposure

[0058] Due to the varying scene brightness, to get a good overall image quality, it is necessary to control the exposure of the CCD to maximize the dynamic range of the digitized image. The main task of exposure control is to keep the sensor operating in the linear range by controlling the shutter speed, and if possible the aperture of the optical system. Since closing the iris and slowing down the shutter speed compensates each other, there exists a certain parameter range in which the exposure remains unchanged. It is obvious that this can be accomplished only to a certain extent as other constraints as capturing fast moving scenes may be desired by the user.

[0059] Besides trying to keep the sensor operating in the linear range it is desirable to maximize the dynamic range of the ADC and hence the digitized image. This is done by controlling the PGA in the AFE. The processing necessary to obtain the relevant control parameters is performed on the DSP.

Auto Focus

[0060] It is also possible to automatically adjust the lens focus in a DSC through image processing. Similar to Auto Exposure, these auto focus mechanisms operate also in a feed back loop. They perform image processing to detect the quality of lens focus and move the lens motor iteratively till the image comes sharply into focus. Auto focus may rely on edge measurements from the edge enhancement previously described.

Playback

[0061] The preferred embodiment DSCs also provide the ability for the user to view the captured images on LCD screen on the camera or on an external TV monitor. Since the captured images are stored in SDRAM (or on compact flash memory) as JPEG bitstreams, playback mode software is also provided on the DSP. This playback mode software decodes the JPEG bitstream, scales the decoded image to the appropriate spatial resolution, and displays it on the

LCD screen and/or the external TV monitor.

Down-sampling

5 [0062] In the preferred embodiment DSC system the image during the playback mode after decoding the JPEG data is at the resolution of the CCD sensor, e.g. 2 Megapixels (1600 x 1200). This image can even be larger depending on the resolution of the CCD sensor. However, for the display purposes, this decoded data has to be down-sampled to NTSC resolution (720 x 480) before it can be fed into the NTSC encoder. Hence, the DSC should implement a down-sampling filter at the tail end of the playback mode thereby requiring additional DSP computation.

10 [0063] The preferred embodiment solves this problem of additional DSP computations by a DCT-domain down-sampling scheme that is included as part of the JPEG decompression module. Note that the JPEG decompression essentially involves three stages: first an entropy decoding stage, followed by an inverse quantization stage, and finally an IDCT stage. In JPEG the IDCT is performed on a block of 8 x 8 pixels. The preferred embodiments down sample a 2 Megapixel image to NTSC resolution (a 4/8 down-sampling) in the IDCT domain by employing a 4 x 4 IDCT to the top left 4 x 4 DCT coefficients (out of a 8 x 8 DCT coefficient block) and hence effectively achieving both the IDCT and the 4/8 down-sampling in one step. The sampling ratio can be varied between 1/8 (smallest image) to 8/8 (full resolution image).

[0064] A separable two-dimensional 4-point IDCT is applied to obtain a 4 x 4 block of image pixels from the top-left (low spatial frequency) 4 x 4 DCT coefficients. By this low-order IDCT we effectively combine anti-aliasing filtering and 8-to-4 decimation. The employed anti-aliasing filter corresponds to a simple operation of preserving only the 16 lowest frequency components in the DCT domain without scaling the preserved DCT coefficients. Though this simple filter is effective in reducing aliasing effect, the preferred embodiments may have a lowpass filter with better frequency response to further reduce aliasing. The use of other lowpass filters will lead to scaling of the preserved coefficients where the scaling factor is the location of each DCT coefficient.

25 [0065] Note that the DCT domain down-sampling technique does not increase the computational complexity. In fact, it reduces the computation since the JPEG decoding stages after entropy decoding does not need to deal with the whole 8 x 8 DCT coefficients except the top-left 4 x 4 coefficients. Use of other anti-aliasing filters also does not add any complexity since the coefficient scaling operation can be merged into the low-order IDCT operation. Also note that this DCT domain down-sampling idea technique can offer n/8 down-sampling ratios, $n = 1, \dots, 7$, for other CCD sensor resolutions.

Up-Sampling

35 [0066] Displaying cropped images for zooming of images also uses an up-sampling scheme. The inverse approach to the down-sampling provides an elegant tool. In the first case the 8x8 DCT coefficients are (virtually) vertically and horizontally extended with zeroes to form a block of NxM coefficients ($N, M > 8$). On this block an IDCT of size NxM is executed yielding NxM samples in the spatial domain.

[0067] Currently, most image pipeline operations are nonstandardized. Having a programmable DSC engine offers the ability to upgrade the software to conform to new standards or improve image pipeline quality. Unused performance can be dedicated to other tasks, such as human interface, voice annotation, audio recording/compression, modem, wireless communication, etc.

40 [0068] Figure 27 shows a preprocessing functional block diagram including CFA interpolation, white balance, color correction, tone scaling, gamma correction, conversion of RGB to YCrCb, edge enhancement, edge detection, color boost, and false color suppression in preparation of JPEG compression. The following sections describe preferred embodiments relating to CFA interpolations.

CFA interpolation with reduced aliasing

50 [0069] A preferred embodiment CFA interpolation for a Bayer pattern (Figure 7a) uses the high-frequency from the green channel to modify the red and blue channel interpolations to reduce the aliasing components at edges within the image by utilizing the signal of the other color channels. By this means artifacts are reduced, sharpness improved, and additional postprocessing avoided. Indeed, proceed as follows.

- (1) apply interpolation to green channel (any interpolation method); this yields the green plane.
- 55 (2) detect edges in the green channel (by gradient or other method).
- (3) compute high-pass component of the green channel (filter with any high-pass filter).
- (4) apply interpolation to the red channel (any interpolation method); this yields the red plane.
- (5) add high-pass component of (3) (with a weighting factor) to red channel.

(6) apply interpolation to the blue channel (any interpolation method); this yields the blue plane.

(7) add high-pass component of (3) (with a weighting factor) to the blue channel.

[0070] So the final image consists of three color planes: the green plane from step (1), the red plane from step (5), and the blue plane from step (7). That is, for a pixel in the final image the green intensity is taken to be the value of the corresponding pixel of the green plane from step (3), the red intensity is taken to be the value of the corresponding pixel of the modified red plane from step (5), and the blue intensity is taken to be the value of the corresponding pixel of the modified blue plane from step (7).

[0071] Theoretical analysis of the foregoing: Each CCD pixel averages the incident optical signal over the spatial extent of the pixel; thus the CCD effectively provides a low-pass filtering of the incident optical signal with a cutoff frequency the reciprocal of the pixel size. Further, the subsampling of the pixel array by the color filters on the pixels leads to aliasing in each color plane. Indeed, for red and blue the subsampling is by a factor of 2 in each direction; so the frequency spectrum folds at half the maximum frequency in each direction. Thus the red and blue baseband spectra areas are each one-quarter of the original array spectrum area (reflecting that the red and blue samplings are each one-quarter of the original array). For green the subsampling is only half as bad in that the spectrum folding is in the diagonal directions and at a distance $\sqrt{2}$ as large as for the red and blue. The green baseband spectrum is one-half the area of the original array spectrum.

[0072] Color fringing at edges is an aliasing problem. In addition, dissimilar baseband spectra lead to color fringing as well, even if no aliasing is present. Indeed, aliasing is not necessarily visible in a single color band image, but the effect becomes obvious upon combination of the three color components into one color image. The shift of the sampling grids between red, green, and blue causes a phase shift of the aliasing signal components. A one-dimensional example clarifies this: presume a one-dimensional discrete signal $f(n)$ and two subsamplings, each by a factor of 2 but one of even-numbered samples and one of odd-numbered samples (so there is a shift of the sampling grids by one sample):

$$f_{\text{even}}(2m) = f(2m)$$

$$f_{\text{even}}(2m+1) = 0$$

$$f_{\text{odd}}(2m) = 0$$

$$f_{\text{odd}}(2m+1) = f(2m+1)$$

Of course, $f(n) = f_{\text{even}}(n) + f_{\text{odd}}(n)$. Let $F(z)$ be the z-transform of $f(n)$, $F_{\text{even}}(z)$ the z-transform of $f_{\text{even}}(n)$, and $F_{\text{odd}}(z)$ the z-transform of $f_{\text{odd}}(n)$. Then noting that $F_{\text{even}}(z)$ is an even function of z (only even powers of z) and $F_{\text{odd}}(z)$ an odd function of z (only odd powers of z):

$$F_{\text{even}}(z) = \{F(z) + F(-z)\}/2$$

$$F_{\text{odd}}(z) = \{F(z) - F(-z)\}/2$$

The $F(-z)$ corresponds to the aliasing and appears with opposite signs; that is, a phase shift of π .

[0073] The color fringing can be reduced by a phase shift of π of the aliased components. However, this is very difficult to achieve, because the only available signal is the sum of the original signal with the aliasing signal. Therefore, the preferred embodiments have another approach.

[0074] As long as two (or more) subsampled signals (i.e., red, green, and blue) have identical characteristics (such as for a gray scale image), a perfect reconstruction of the original image can be achieved by just adding the subsampled signals. However, in CFA interpolation generally the subsampled signals stem from different color bands. Aliasing errors become visible especially at edges where the interpolated signals of the different color bands are misaligned. Therefore, the preferred embodiments counter color fringing at edges by reducing the aliasing components only at edges through utilization of other ones of the subsampled signals. This reduces artifacts, improves sharpness, and avoids additional postprocessing.

[0075] In particular, for Bayer pattern CFA the green channel has a higher cutoff frequency than that of the red and

blue channels; thus the green channel has less significant aliasing. The aliasing signal to be compensated is a high-pass signal, which is now estimated as the high-pass component of the green channel; and this is added (rather than subtracted due to the phase shift due to the offset of the red and blue subsampling grids relative to the green subsampling grid) to the red and blue channels. The high-pass green component could be multiplied by a scale factor prior to addition to the red and blue subsamplings. The signals are added while interpolating red, blue or afterwards.

CFA interpolation with inter-hue adaptation

[0076] Alternative CFA interpolation preferred embodiments first interpolate Bayer pattern greens using a 5x5 FIR filter, and then use the interpolated green to interpolate red and blue each with two steps: first interpolate diagonally to form a pattern analogous to the original green pattern (this interpolation uses a normalization by the green to estimate high frequencies), and then apply a four-nearest neighbor interpolation (again using green normalization to estimate high frequencies) to complete the red or blue plane.

[0077] More explicitly, denote the CFA value for pixel location (y,x) , where y is the row number and x the column number of the array, as follows: red values $R(y,x)$ at pixel locations (y,x) where y and x are both even integers, blue values $B(y,x)$ where y and x are both odd integers, and green values $g(y,x)$ elsewhere, that is, where $y+x$ is an odd integer.

[0078] First, let $G^\wedge(y,x)$ denote the green value at pixel location (y,x) resulting from the green plane interpolation; this is defined for all pixel locations (y,x) . This interpolation can be done by various methods, including the edge preservation interpolation of the following section. Note that many interpolations do not change the original green values; that is, $G^\wedge(y,x) = G(y,x)$ may be true for (y,x) where G was originally defined (i.e., $y+x$ is an odd integer).

[0079] Next, define the red and blue interpolations each in two steps as illustrated in Figure 28 which is labeled for blue and uses arrows to show interpolation contributions.

[0080] First red step: $R(y,x)$ is already defined for pixel locations (y,x) with $y=2m$, and $x=2n$ with m and n integers; so first for $y=2m+1$ and $x=2n+1$, define $R^\wedge(y,x)$:

$$R^\wedge(y,x) = G^\wedge(y,x) \{ R(y-1,x-1) / G^\wedge(y-1,x-1) + R(y-1,x+1) / G^\wedge(y-1,x+1) + \\ R(y+1,x-1) / G^\wedge(y+1,x-1) + R(y+1,x+1) / G^\wedge(y+1,x+1) \} / 4$$

This interpolates the red plane to the pixels where $B(y,x)$ was defined. (Figure 28 illustrates the analogous interpolation for blue.) Note that this interpolation essentially averages the red values at the four corners of the 3x3 square about (y,x) with the values normalized at each location by the corresponding green values. If any of the green values are below a threshold, then omit the normalization and just average the red values.

[0081] Perform the first blue step in parallel with the first red step because the same green values are being used.

[0082] First blue step: $B(y,x)$ is already defined for pixel locations (y,x) with $y=2m+1$, and $x=2n+1$ with m and n integers, so first for $y=2m$ and $x=2n$, define $B^\wedge(y,x)$:

$$B^\wedge(y,x) = G^\wedge(y,x) \{ B(y-1,x-1) / G^\wedge(y-1,x-1) + B(y-1,x+1) / G^\wedge(y-1,x+1) + \\ B(y+1,x-1) / G^\wedge(y+1,x-1) + B(y+1,x+1) / G^\wedge(y+1,x+1) \} / 4$$

This interpolates the blue plane to the pixels where $R(y,x)$ was defined as illustrated in the lefthand portion of Figure 28. Again, this interpolation essentially averages the blue values at the four corners of the 3x3 square about (y,x) with the values normalized at each location by the corresponding green values.

[0083] Second red step: define $R^\wedge(y,x)$ where $y+x$ is an odd integer (either $y=2m$ and $x=2n+1$ or $y=2m+1$ and $x=2n$)

$$R^\wedge(y,x) = G^\wedge(y,x) \{ R^\wedge(y-1,x) / G^\wedge(y-1,x) + R^\wedge(y,x+1) / G^\wedge(y,x+1) + R^\wedge(y+1,x) / G^\wedge(y+1,x) + R^\wedge(y,x-1) / G^\wedge(y,x-1) \} / 4$$

This second step interpolates the red plane portion defined by the first step to the pixels where $G(y,x)$ is defined. Again, this interpolation essentially averages the red values at four neighboring pixels of (y,x) with the values normalized at each location by the corresponding green values.

[0084] Second blue step: define for $y+x$ an odd integer (either $y=2m$ and $x=2n+1$ or $y=2m+1$ and $x=2n$)

$$B^{\wedge}(y,x) = G^{\wedge}(y,x) \{ B^{\wedge}(y-1,x) / G^{\wedge}(y-1,x) + B^{\wedge}(y,x+1) / G^{\wedge}(y,x+1) + B^{\wedge}(y+1,x) / G^{\wedge}(y+1,x) + B^{\wedge}(y,x+1) / G^{\wedge}(y,x+1) \} / 4$$

5

This second step interpolates the blue plane portion defined by the first step to the pixels where $G(y,x)$ is defined. Again, this interpolation essentially averages the blue values at four neighboring pixels of (y,x) with the values normalized at each location by the corresponding green values.

[0085] The final color image is defined by the three interpolated color planes: $G^{\wedge}(y,x)$, $R^{\wedge}(y,x)$, and $B^{\wedge}(y,x)$. The particular interpolation used for $G^{\wedge}(y,x)$ will be reflected in the normalizations for the two-step interpolations used for $R^{\wedge}(y,x)$ and $B^{\wedge}(y,x)$.

CFA interpolation with edge preservation

[0086] Alternative CFA interpolation preferred embodiments interpolate Bayer pattern greens by a (small) FIR filter plus preserve edges by a comparison of an interpolated pixel green value with the nearest-neighbor pixel green values and a replacement of the interpolated value with a neighbor value if the interpolated value is out of range. Figure 29 illustrates the green interpolation. After this green interpolation, interpolate the red and blue planes.

[0087] In particular, first at each pixel (y,x) apply the following 5x5 FIR filter to $G(y,x)$ defined on the pixels (y,x) where $x+y$ is odd to yield $G1(y,x)$ defined for all (y,x) :

$$1/200 \begin{vmatrix} 0 & -11 & 0 & -11 & 0 \\ -11 & 0 & 72 & 0 & -11 \\ 0 & 72 & 200 & 72 & 0 \\ -11 & 0 & 72 & 0 & -11 \\ 0 & -11 & 0 & -11 & 0 \end{vmatrix}$$

30

The 200 center entry just implies for (y,x) where $G(y,x)$ is defined in the CFA, $G1(y,x) = G(y,x)$. Note that green values are in the range of 0-255, and negative values are truncated to 0. Of course, other FIR filters could be used, but this one is simple and effective.

[0088] Next, for the (y,x) where $G1(y,x)$ is interpolated, consider the four nearest neighbors' values $G(y\pm 1,x)$, $G(y,x\pm 1)$ and discard the largest and smallest values. Let A and B be the remaining two nearest-neighbor values with B greater than or equal to A. Then define the final interpolated green value $G^{\wedge}(y,x)$ as follows:

$$G^{\wedge}(y,x) = \begin{cases} A & \text{if } G1(y,x) < A \\ G1(y,x) & \text{if } A \leq G1(y,x) \leq B \\ B & \text{if } B < G1(y,x) \end{cases}$$

45

This clamps the interpolated value to midrange of the neighboring pixel values and prevents a single beyond-the-edge nearest-neighbor pixel from diluting the interpolated pixel value. Figure 29 shows the overall green interpolation.

[0089] Complete the image by red and blue interpolations. The red and blue interpolations may each be a single step interpolation, or each be a two-step interpolation as described in the foregoing section which uses the edge-preserved green values, or each be some other type of interpolation.

CFA interpolation plus noise filtering

[0090] Preferred embodiments save on line memory required for CFA interpolation followed by lowpass filtering to limit noise with an integrated approach. In particular, CFA interpolation typically contains a horizontal interpolation block and a vertical interpolation block with line memories in between as illustrated in Figure 30. The horizontal interpolation block has an input of a row of CFA signals, two toggle switches, two zero insertion subblocks, two three-tap FIR filters (coefficients 0.5, 1.0, 0.5), and two outputs: one output for each color. Each of the FIR filters just reproduces the input

color values and puts the average of successive input color values in place of the inserted zeros. The zero-insertion and toggle timing of two subblocks alternate with each other. The block diagram of the horizontal interpolation block is shown in Figure 31 with a row of raw data R/G/R/G/R ... : in this block row-interpolated Red and Green signals are output. In case the row of raw data input is B/G/B/G/B... interpolated Blue and Green signals are output.

[0091] A line (row) memory delays the data by one CFA line (row) period in order to interpolate the data in the vertical interpolation block. Figure 32 shows the four line memories and the input/output data of the memories. In the case of an input row of R/G/R/G/... raw data with m indicating the (even) row number and n the column number which increments as the row data enters, the input and output data are:

10 Input_A = R(m,n)
 Output_A1 = Input_A = R(m,n)
 Output_A2 = G(m-1,n) which was the interpolated green from the previous row of raw data, a G/B/G/B... row
 Output_A3 = R(m-2,n) which was the interpolated red from the second previous row of raw data, a R/G/R/G/... row
 Input_B = G(m,n)
 15 Output_B1 = Input_B = G(m,n)
 Output_B2 = B(m-1,n) which was the interpolated blue from the previous row of raw data, a G/B/G/B/... row
 Output_B3 = G(m-2,n) which was the interpolated green from the second previous row of raw data, a R/G/R/G/... row

20 **[0092]** This provides the two rows of red, R(m,n) and R(m-2,n), for vertical interpolation to create the m-1 row of red and also provides the green rows G(m,n), G(m-1,n), and G(m-2,n) which do not need vertical interpolation.

[0093] The next input row (row m+1) of G/B/G/B/... raw data leads to the following input and output data:

 Input_A = G(m+1,n)
 25 Output_A1 = Input_A = G(m+1,n)
 Output_A2 = R(m,n) which was the interpolated red from the previous row of raw data, a R/G/R/G/... row
 Output_A3 = G(m-1,n) which was the interpolated green from the second previous row of raw data, a G/B/G/B/... row
 Input_B = B(m+1,n)
 30 Output_B1 = Input_B = B(m+1,n)
 Output_B2 = G(m,n) which was the interpolated green from the previous row of raw data, a R/G/R/G/... row
 Output_B3 = B(m-1,n) which was the interpolated blue from the second previous row of raw data, a G/B/G/B/... row

 This provides the two rows of blue, B(m+1,n) and B(m-1,n), for vertical interpolation to define the m row blue and also provides the green rows G(m+1,n), G(m,n), and G(m-1,n) which do not need vertical interpolation.

[0094] Figure 33 shows the combinations for vertical interpolations. In particular, for row m output (row m+1 input) the combinations are (Figure 33b):

40 green is G(m,n)
 red is R(m,n)
 blue is (B(m-1,n) + B(m+1,n))/2

And for row m-1 output (row m input) the combinations are (Figure 33a):

45 green is G(m-1,n)
 red is (R(m,n) + R(m-2,n))/2
 blue is B(m-1,n)

50 **[0095]** As Figure 33 illustrates, a vertical low-pass noise filter can be applied directly to the three green outputs (G(m-2,n), G(m-1,n), and G(m,n) for row m input and G(m-1,n), G(m,n), and G(m+1,n) for row m+1 input), but red and blue cannot be vertically filtered because the four line memories of Figure 32 do not output enough lines (rows). Rather, eight line memories are needed as illustrated in Figure 34.

55 **[0096]** Figures 35a-35b illustrate the preferred embodiment combination vertical interpolation and low-pass noise filtering including green vertical noise reduction filter block A, green-noise block B, blue/red green-noise difference block C, and red/blue green-noise sum block D. The six inputs for the preferred embodiments of Figures 35a-35b are the outputs of the horizontal interpolations and four line memories of Figures 30-32 and thus the same as the inputs to the known vertical interpolation filter of Figure 34.

[0097] For an implementation of this interpolation plus noise filtering on a programmable processor the eight line

memories in Figure 34 would take up twice as much processor memory space as the four line memories of Figures 30-32, and this can be significant memory space. For a large CFA such as a 2 megapixel (1920 by 1080 pixels) CCD, a line memory would be 1-2 kbytes, so the difference would be 4-8 kbytes of processor memory.

[0098] In more detail, Figure 35a illustrates the noise reduction and vertical interpolation for the case of input row m with m an even integer (raw CFA data $R/G/R/G/\dots$) into the horizontal interpolator plus four line memories of Figure 32: the six (horizontally interpolated) inputs at the lefthand edge of Figure 35a are $R(m,n)$, $G(m-1,n)$, $R(m-2,n)$, $G(m,n)$, $B(m-1,n)$, and $G(m-2,n)$ (i.e., the outputs in Figure 32); and the output will be noise-reduced colors for row $m-1$: $R''(m-1,n)$, $G''(m-1,n)$, and $B''(m-1,n)$. First, the vertical interpolation (lefthand portion of Figure 35a) averages $R(m,n)$ and $R(m-2,n)$ to create $R(m-1,n)$; $G(m-1,n)$ and $B(m-1,n)$ already exist as inputs.

[0099] Then the noise reduction filter (block A in the righthand portion of Figure 35a) creates and outputs the vertically low-pass filtered green $G''(m-1,n)$ as:

$$G''(m-1,n) = [G(m,n) + 2 \cdot G(m-1,n) + G(m-2,n)]/4$$

[0100] Next, block B creates Δ_G as the difference between G and G'' ; that is, Δ_G is the vertical high-frequency part of G :

$$\Delta_G(m-1,n) = G(m-1,n) - G''(m-1,n)$$

Because G is sampled twice as frequently as B and R in the Bayer CFA, direct high-frequency estimation of G will likely be better than that of B and R , and thus the preferred embodiment uses Δ_G to subtract for noise reduction. Note that the difference between the vertical average $[G(m+1,n) + G(m-1,n)]/2$ and $G''(m,n)$ equals $-\Delta_G(m,n)$, so for R and B which are to be vertically interpolated (averaged) plus low-pass filtered, the high-frequency estimate provided by G which is to be subtracted from R and B will have opposite sign.

[0101] Thus block C subtracts Δ_G from B to create B'' for row $m-1$ because B is not vertically interpolated for $m-1$:

$$B''(m-1,n) = B(m-1,n) - \Delta_G(m-1,n)$$

Essentially, the vertical high-frequency part of G is used as an estimate for the vertical high-frequency part of B , and no direct vertical low-pass filtering of B is applied.

[0102] Then block D adds Δ_G to R to create R'' for row $m-1$ because R was vertically interpolated:

$$R''(m-1,n) = R(m-1,n) + \Delta_G(m-1,n)$$

Again, the vertical high-frequency part of G is used in lieu of the high-frequency part of R , and because an vertical averaging creates $R(m-1,n)$, the opposite sign of Δ_G is used to subtract the high-frequency estimate.

[0103] Thus the noise-reduced filtered three color output row $m-1$ are the foregoing $G''(m-1,n)$, $R''(m-1,n)$, and $B''(m-1,n)$.

[0104] Similarly, for output row m from input row $m+1$ (again with m an even integer) and raw CFA data $G/B/G/B/\dots$ the six (horizontally interpolated) inputs are $G(m+1,n)$, $R(m,n)$, $G(m-1,n)$, $B(m+1,n)$, $G(m,n)$, and $B(m-1,n)$, and the output will be noise-reduced colors for row m : $R''(m,n)$, $G''(m,n)$, and $B''(m,n)$. The vertical interpolation (lefthand portion of Figure 35b) averages $B(m+1,n)$ and $B(m-1,n)$ to create $B(m,n)$; $G(m,n)$ and $R(m,n)$ already exist as inputs. Then the noise reduction filter (righthand portion of Figure 35b) block A again creates vertically low-pass filtered green $G''(m,n)$ as:

$$G''(m,n) = [G(m+1,n) + 2 \cdot G(m,n) + G(m-1,n)]/4$$

Next, block B again creates the vertical high-frequency portion of G , called Δ_G , as the difference between G and G'' :

$$\Delta_G(m,n) = G(m,n) - G''(m,n)$$

[0105] Then block C again subtracts Δ_G but from R (rather than B as for row $m-1$ outputs) to create R'' :

$$R''(m,n) = R(m,n) - \Delta G(m,n)$$

Thus the high-frequency part of G is again used as an estimate for the noisy part of R, and no direct noise filtering of R is applied, but for row m the Delta_G is subtracted rather than added as for row m-1. Indeed, for R even rows have Delta_G subtracted and odd rows have Delta_G added because the odd rows have R defined as a vertical average.
 [0106] Lastly, block D adds Delta_G to B to create B'':

$$B''(m,n) = B(m,n) + \Delta G(m,n)$$

Thus as with R, the Delta_G vertical high-frequency estimate is row-by-row alternately added to and subtracted from B instead of a direct vertical low-pass filtering of B. Note that for a given row the Delta_G terms for R and B have opposite signs because one of R and B will be an average of preceding and succeeding rows.

[0107] In short, the preferred embodiments are able to emulate the CFA horizontal interpolation, vertical interpolation, and low-pass filtering with only four line memories by using a high-frequency estimate based on G.

[0108] Figures 36a-36b and 37a-37b illustrate an alternative embodiment in which the vertical low-pass filtering of G differs from the 1/4, 1/2, 1/4 weighting of the preferred embodiments of Figures 35a-35b.

CFA interpolation for complementary color CCD

[0109] Preferred embodiment CFA interpolations for a complementary color pattern CFA (illustrated in Figure 7b) combine a simple interpolation followed by an image quality enhancement by detection and adjustment for color imbalance. In particular, presume initial interpolation as defined at each pixel all four complementary color values, and denote the color values as Ye (yellow), Cy (cyan), Mg (magenta), and G (green).

[0110] First, at each pixel compute an imbalance factor μ :

$$\mu = Ye + Cy - 2 \cdot G - Mg$$

This imbalance factor represents the difference between ideal and actual pixel color values. Indeed, the definitions of the complementary color values in terms of red value (R), green value (G), and blue value (B) are $Ye = R + G$, $Cy = G + B$, and $Mg = B + G$. Hence, the following relation always holds for a pixel's color values:

$$Ye + Cy = 2 \cdot G + Mg$$

Thus the imbalance factor μ ideally vanishes. When an edge is near a pixel, imbalance can arise due to the spatial difference of each of the four color samples in the CFA. The preferred embodiments detect the imbalance and adjust by modifying each color value:

$$Ye' = Ye - \mu/4$$

$$Cy' = Cy - \mu/4$$

$$Mg' = Mg + \mu/4$$

$$G' = G + \mu/8$$

Then these modified complementary colors are used to form the final image.

[0111] Figure 38 illustrates the overall flow for the enhancement using the imbalance factor. Of course, scale factors other than -1/4, -1/4, 1/4, and 1/8 could be applied to the imbalance factor provided that $Ye' + Cy' = 2 \cdot G' + Mg'$.

White balance

[0112] The term "white balancing" is typically used to describe algorithms, which correct the white point of the camera with respect to the light source under which the camera currently operates. Since the estimation of the true light spectrum is very difficult, the aim of most approaches is to correct the output of the red and blue channel (assuming CCDs based on the RGB color filters), such that for a gray object the pixel intensities for all color channels are almost identical. The most common technique basically calculates the average energy or simply the mean for each channel. The calculation of averages may be carried out in N local windows W_j , $j = 1, 2, \dots, N$, as for red:

$$R_j = \sum_{k \in W_j} r(k)$$

with $r(k)$ denoting the digital signal for the red channel. Similar averages B_j and G_j are calculated for the blue and green color channels. The imbalance between the channels, given by the green-to-red and green-to-blue ratios

$$WBR = \sum_j G_j / \sum_j R_j$$

$$WBB = \sum_j G_j / \sum_j B_j$$

are used as correction multiplier for the red and blue channels, respectively

$$r'(k) = WBR \, r(k)$$

$$b'(k) = WBB \, b(k)$$

There exist many different flavors of this approach, which all calculate intensity-independent multiplication factors WBR and WBB .

[0113] This approach works only if several assumptions are valid. First, it is assumed that the sensor responses are well aligned over the input intensity range; in other words, the green response curve equals the red (blue) response curve multiplied by a factor. Looking at sensor (CCD) characteristics indicates that this assumption does not hold. For high light intensities, the sensor saturates; while at very low light intensities, the sensor response (especially for the blue channel) is very small. Furthermore, non-linearities of the sensor, as well as imbalances of the color channels related to the sensor response and the light source, are handled simultaneously. Resulting artifacts include magenta colors in very bright areas, where the "color" should turn white, or wrong colors in dark areas.

[0114] The pixel intensity at the sensor output, e.g. for the red color channel, can be modeled as

$$r(k) = \int I(\lambda) \beta(k, \lambda) f_R(\lambda) \alpha(I, \lambda) d\lambda$$

where λ denotes the wavelength, $I(\lambda)$ the spectrum of the light source, $\beta(x, \lambda)$ the reflectance of the object under observation, $f_R(\lambda)$ the spectral sensitivity of the red color filter covering the CCD pixels, and $\alpha(I, \lambda)$ the intensity- and wavelength-dependent efficiency of the CCD in converting photons into electrons.

[0115] Regarding only the spectral response curves of the color filters $f_R(\lambda)$ (and also $f_G(\lambda)$ and $f_B(\lambda)$) of a typical CCD sensor, the output signals differ:

$$WBR = \int f_G(\lambda) d\lambda / \int f_R(\lambda) d\lambda = 1.09$$

$$WBB = \int f_G(\lambda) d\lambda / \int f_B(\lambda) d\lambda = 1.34$$

The values are obtained using the response of a typical CCD and assuming perfect white light source (the spectrum $I(\lambda)$ is flat), a perfectly white object (the spectrum of the reflected light is identical to the spectrum of the illuminating light which means $\beta(k, \lambda) = 1$), and neglecting $\alpha(I, \lambda)$ (no wavelength dependent quantum efficiency). Especially the blue

channel shows a smaller response than green or red at the same intensity. The non-linear quantum efficiency of the sensor is another effect. A typical s-shaped sensor response over the input intensity is shown in Figure 39a. Furthermore, the sensor response in each channel depends on spectrum of the light source.

[0116] Thus, preferred embodiment white balancing takes into account the misalignment as well as the non-linearity. Typical light sources are not flat over the visible spectrum but tend to have a higher energy in certain spectral bands. This effect influences the observed sensor response; ideally it should be corrected by white point compensation, which may be based on a correction matrix. An independent balancing of the channels cannot handle this effect as previously outlined. For ease of mathematical description, approximate the s-shaped response curve in Figure 39a by piecewise linear segments. Three segments separate the light conditions into three categories: very low intensity, normal intensity, and very bright light. Figure 39b shows the effect of applying a single multiplier. With respect to the green signal, the amplification of the blue signal is too small in low light conditions, whereas in very bright conditions the multiplier is too large. Reducing the factor leaves an offset between the components, visible as wrong colors. Therefore, the correction terms for aligning all three response curves must look different and reflect the sensor characteristics.

[0117] The preferred embodiment white balancing splits into two separate schemes, one accounts for imager dependent adjustments, while the other one is related to light sources.

[0118] Without any restrictions on generality, the s-shape response curve is approximated in the following by three piecewise linear segments. More segments increase the accuracy but do not change the basic concept. For the first region (very low intensity) and the blue channel, the model reads with s the response and x the input intensity:

$$S_{B,1} = a_{B,1}x$$

Modeling the second region requires a multiplier and an offset

$$S_{B,2} = a_{B,2}x + b_{B,2}$$

The offset term is determined by the constraint that the response curve needs to be contiguous at the transition point x_1 from region 1 to region 2:

$$s_{B,1}(x_1) = s_{B,2}(x_1)$$

$$\text{so } b_{B,2} = (a_{B,1} - a_{B,2})x_1$$

The parameters for the linear model of region 3

$$s_{B,3} = a_{B,3}x + b_{B,3}$$

are completely determined because the maximum output has to be identical to the maximum input x_{\max} , and the response curve needs to be contiguous at the joint point x_2 :

$$x_{\max} = a_{B,3}x_{\max} + b_{B,3}$$

$$s_{B,2}(x_2) = s_{B,3}(x_2)$$

$$a_{B,3} = (s_{B,2}(x_2) - x_{\max}) / (x_2 - x_{\max})$$

$$b_{B,3} = (1 - a_{B,3})x_{\max}$$

Thus the parameters to specify the approximation of the response curve for each color component are a_1 , a_2 , x_1 , and x_2 . x_{\max} is not a free parameter, because it is specified by the bit resolution of the input signal.

[0119] The preferred embodiment white balancing now applies different multipliers for each region. For continuous transition from one region to the next, an additional offset is required. Although the number of regions is arbitrary, without loss of generality only three regions are considered in the following equations. The correction term for blue with respect to green for region 1 has to be:

$$WBB_1 = a_{G,1} / a_{B,1} \approx G_1 / B_1$$

where window 1 (for G_1 and B_1) has pixels with intensities in region 1.

Thus, an input intensity value lying in region 1 gets the corrected output

$$b'(k) = WBB_1 b(k)$$

Based on the balancing multiplier for region 2

$$WBB_2 = a_{G,2} / a_{B,2} \approx G_2 / B_2$$

the white balancing must consider an additional offset for values in region 2

$$b'(k) = WBB_2 b(k) + WBOB_2$$

with

$$WBOB_2 = (WBB_1 - WBB_2) x_1$$

For the third region the calculation is basically the same, except that no explicit WBB_3 can be specified, but the amplification is determined by the maximum value x_{max} .

$$b'(k) = WBB_3 b(k) + WBOB_3$$

with

$$WBB_3 = (x_{max} - (WBB_2 x_2 + WBOB_2)) / (x_{max} - x_2)$$

$$WBOB_3 = (1 - a_{B,3}) x_{max}$$

For an implementation, the system must determine appropriate white balancing multipliers WBB_i for $N-1$ regions. Based on these values, the remaining offset values $WBOB$ and the multiplier for the last regions are calculated. The locations of the transition points are specified a priori. The white balancing itself selects the region based on the intensity value of the input pixel and applies the appropriate gain and offset to that value

$$b'(k) = \begin{cases} WBB_1 * b(k) & b(k) \leq x_1 \\ WBB_2 * b(k) + WBOB_2 & x_1 < b(k) \leq x_2 \\ WBB_3 * b(k) + WBOB_3 & x_2 < b(k) \end{cases}$$

Plus a similar multiplier for the red channel.

[0120] The total dynamic range of the CCD output signal is independent of aperture, and shutter, since they affect

the number of photons captured in the CCD. An analog gain however, or any digital gain prior to processing shifts the signal and should be avoided. In case a gain (digital) α needs to be applied, this gain can be included into the white balancing method. A gain maps the maximum input value x_{\max} to the output value $\alpha \cdot x_{\max}$.

[0121] The scaled response curves behave identical to the non-scaled one, meaning that the scaled signal saturates at $\alpha \cdot x_{\max}$. Substituting

$$WBB_1 := \alpha \cdot WBB_1$$

$$WBB_2 := \alpha \cdot WBB_2$$

In that way the equation in the previous section remain unchanged, except

$$WBOB_3 = (\alpha - a_{B,3}) x_{\max}$$

After linearization the signal can undergo an adjustment reflecting the light source. This is also known as white point adjustment. Here the input signal is transformed such that it looks like as if it has been captured under a different light source. For example, an image has been captured in bright sunlight (D65), but the color characteristics should be as if it has been captured under indoor conditions (D50 tungsten).

$$[R, G, B] D_{65}^T = I_{D65}^T \cdot \beta \cdot [f_R, f_G, f_B]^T$$

$$[R, G, B] D_{50}^T = I_{D50}^T \cdot \beta \cdot [f_R, f_G, f_B]^T$$

Here, I_{Dxx} denotes a vector sampling the light spectrum, β is a diagonal matrix describing the reflectance of the objects, and f_R , f_G , and f_B denote the spectral response of the CCD light filters. Based on these equations a 3x3 transformation matrix can be calculated relating the signal under D65 to D50:

$$[R, G, B] D_{50}^T = I_{D50}^T \cdot I_{D65}^{-T} \cdot [R, G, B] D_{65}^T$$

The 3x3 transformation matrix

$$M_D = I_{D50}^T \cdot I_{D65}^{-T}$$

can be calculated offline.

In real systems it is almost impossible to determine averages for the different response regions. Therefore a simple solution is to calculate overall values as in the foregoing ratio of integrals, and modify them with fixed values based on predetermined sensor measurements

$$WBB_1 = \alpha_1 \cdot WBB$$

$$WBB_2 = \alpha_2 \cdot WBB$$

And similarly for WBR.

The transition points can be fixed in advance, too. There is just one exception for the transition point x_2 . In rare situations the WBR-value may be so large that it exceeds the maximum output value at the transition point x_2 . In that situation, either the WBR needs to be decreased or the transition point is reduced. The diagram in Figure 40 shows an example of the effectiveness of this technique. The red components is adjusted with respect to the green component. Using a single multiplier exceeds the green signal in bright areas, and is less effective in low light areas, whereas the segmented white balancing matches the green curve for all intensities.

Resizing preferred embodiments

[0122] Frequently images captured in one size (e.g., 320 x 240 pixels) have to be converted to another size (e.g., about 288 x 216) to match various storage or input/output formats. In general this requires a fractional up-sampling or down-sampling by a rational factor, N/M ; for example, a resizing from 320 x 240 to 288 x 216 would be a 9/10 resizing. Theoretically, resizing amounts to cascaded interpolation by N , anti-aliasing filter, and decimation by M . In practice the resizing may be achieved with an M -phase, K -tap filtering plus selection of N outputs per M inputs.

[0123] For example, preliminarily consider a resizing by a ratio of 63/64 using a 3-tap filter as illustrated in Figure 41a in which the top horizontal line represents pixel inputs and the horizontal length-three braces represent the 3-tap filter kernel applied to the indicated three inputs and producing the indicated outputs. Indeed, presume the filter kernel is a continuous function $f(t)$ with support of length 3-1/63 so that at most three inputs can be involved; see Figure 41b. Note the slight shifting to the right of successive braces in Figure 41a: this represents the resizing from 64 inputs down to 63 outputs because the center of the filter kernel (and thus the non-rounded-off output position) must increment $1+1/63$ (=64/63) pixel positions for each output in order for the 63 outputs to match the 64 inputs. Output[0] (represented by the farthest left brace in Figure 41a) is centered at the position of input, and the non-rounded-off output position j , denoted $\text{outp_pos}[j]$, thus equals $1 + j \cdot 64/63$.

[0124] The filter kernel is represented as a symmetrical continuous function $f(t)$ centered at time 0. Output[0] for example, needs three kernel values: $f(-1)$, $f(0)$, and $f(1)$. Each output point is computed as the inner product of three kernel coefficient values with three input pixel values. The center input point for the output[j] is positioned at $\text{round}(\text{outp_pos}[j])$ where $\text{round}()$ is the round off function. The other two input points are offset from this center point by ± 1 . The center filter kernel coefficient value is $f(\text{round}(\text{outp_pos}[j]) - \text{outp_pos}[j])$ and the other are $f()$ at the ± 1 offsets of this center value point. Thus the following table shows the output position, coefficient kernel values, and input points needed for each output:

output j	outp_pos	center coeff position	input points
0	1	0	0,1,2
1	2 1/63	-1/63	1,2,3
2	3 2/63	-2/63	2,3,4
...
31	32 31/63	-31/63	31,32,33
32	33 32/63	31/63	33,34,35
33	34 33/63	30/63	34,35,36
...
61	62 61/63	2/63	62,63,64
62	63 62/63	1/63	63,64,65
63	65	0	64,65,66
...

The table shows the desired coefficient position as well as the inputs involved in each output. Note the $j=63$ case is similar to the $j=0$ case in that the kernel center aligns with the input, but with the output position and input indices shifted by 64. Notice that at $j=32$ there is a change in the input pattern: for $j \leq 31$, output[j] uses input j , $j+1$, and $j+2$; whereas for $j \geq 32$, output[j] uses inputs $j+1$, $j+2$, and $j+3$.

[0125] The preferred embodiments partition the filtering computations for resizing a two-dimensional array (image) between iMX 124 and DSP 122 and limit memory use as follows. First iMX 124 performs the 3-tap row filtering with 64 banks of coefficients and then 3-tap column filtering with 64 banks of coefficients. First consider the row filtering. 3-tap row filtering on iMX 124 has the input/output relationship:

iMX output j	input points
0	0,1,2
1	1,2,3

(continued)

iMX output j	input points
2	2,3,4
...	...
31	31,32,33
32	32,33,34
33	33,34,35
...	...
61	61,62,63
62	62,63,64
63	63,64,65
64	64,65,66
...	...

Comparing this table with the prior 63/64 resizing table shows that the only difference is the iMX produces one extra point, namely, IPP_output[32]. Thus the preferred embodiments produce the 64 output points with iMX 124, and then use DSP 122 to pick the 63 valid points:

output [j] = IPP_output[j]

for j = 1,2,...,31 IPP_output[j+1] for j = 32,33, ..., 62

In general, N/M resizing when N/M is less than 1 involves deleting M-N outputs of every M outputs. Thus the preferred embodiments generally perform the filter operations on the M input points in an accelerator such as the iMX and then use a processor such as the DSP to discard the unneeded outputs. (iMX can also handle larger-than-unity resizing up to N/M = 3.)

[0126] iMX can produce 8 outputs of 3-tap row filter in 3 cycles. Basically, 8 adjacent outputs are computed in parallel using the 8 MAC units. At time 0, pull out input points 0,1,2,3,...,7, multiply with appropriate coefficients (each can be different), and accumulate into 8 accumulators. At time 1 pull out input points 1,2,...,8, do the same, and at time 2, pull out input points 2,3,...,9, accumulate the products, and write out 8 outputs, j=0,1,...,7. Next, shift over 8 input points to compute j=8,9,...,15.

[0127] For the vertical direction, iMX computes 8 outputs in parallel as well. These are 8 horizontally adjacent output points, and every fetch of input array also bundles 8 horizontally adjacent output points. Therefore, all 8 MAC units share the same coefficient values for each cycle. For vertical direction there is less data reuse in iMX, so input/output memory conflicts slow down the computation to 4 cycles/8 outputs. Total filtering time is 7 cycles/8 outputs, or 7/8 cycle per output.

Input data is of size 320x240x3. Thus, the filtering of iMX takes 320x240x3.7/8 201,600 cycles, or 1.7 msec with iMX running at 120 MHz.

[0128] After filtering, DSP picks correct outputs. Basically, one row out of every 64 rows and one column out of every 64 columns should be discarded. A DSP assembly loop moves the valid iMX output points to a separate output area. iMX and DSP may run in parallel if there is sufficient local memory for both. An entire input image likely is too large to fit into local memory; even the natural choice, 63x63 output points, may be too large. In such a case partition the image, such as 63 wide x 16 tall, and deal with extra bookkeeping in the vertical direction. With just 3x64 = 192 coefficients, it would be economical to pre-compute and store them. DSP should keep track of the phase of each processing block, and point iMX to the correct starting address of coefficients. If the colors are interleaved, this allows interleaved filtering as well. iMX deals with strides in-getting input points. The following table shows interleaved 3-tap filtering.

j	input points
0	0,3,6
1	1,4,7
2	2,5,8
...	...

However, interleaving consumes three times more memory for the same output block size for each color. Thus it is possible to partition the task into smaller size, such as 63×5 on each color plane, and deal with extra overhead in the vertical direction. If the color format is not 4:4:4 (say, 4:2:2), and input is colorinterleaved, the DSP will need to spend some additional time separating color planes.

[0129] Performing resizing totally in DSP 122 is time-consuming if implemented with straightforward fractional addressing. The preferred embodiments streamline the computation by requiring filter coefficients to be reordered and padded with dummy words. iMX 124 performs the main processing concurrently with DSP 122 computing the coefficients. This efficiently realizes high throughput resizing.

[0130] In more detail, the preferred embodiments perform an N/M resizing of an image by using iMX 124 to perform M-phase, K-tap filtering (which produces redundant output points) and DSP 122 to select the correct output points. Further, DSP 122 computes needed coefficients from a fewer-subsample coefficient template to reduce memory usage to 8^*K ; otherwise memory usage up to 2^*M^*K coefficient words would be needed. DSP 122 can compute the rounded position for the coefficients, and build up the coefficient memory for iMX 124.

[0131] For processing wide and short blocks of pixels (i.e., 16×64) the horizontal direction requires more computation in that horizontal coefficients are updated more often than vertical coefficients. However, the coefficients constructed by DSP 122 can be reused many times within the short block, so the load on DSP 122 should not be excessive.

[0132] In particular, preferred embodiments proceed with the following steps which are illustrated in Figures 42a-42e for a 3tap filter and a 10-to-9 resizing (e.g., resizing from 320×240 to 288×216 in 30 frames/sec) (presume 4:4:4 interleaved, for 4:2:2 or 4:1:1 do subsampling after resizing):

1. select input/output pattern: every 10 inputs leads to 9 outputs as per Figure 42a.
2. draw coefficient pattern for a processing unit, one color first. Arrows in Figure 42b indicate which input points are used: connected arrows form the same output point, and gray (open head) arrows indicate zero coefficients. Thus three input points determine the first output point, only two input points determine each of the next eight output points, and then a tenth ignored output (no non-zero input points); and this repeats every ten. This pattern suggests use of a polyphase 3-tap filter, and drop the last output in every group of 10 outputs.
3. consider interleaved input/output. See Figure 42c which shows a set of three groups of ten input points interleaved so that the three input points determining the first output point from the original first group of ten input points are now at locations 1, 4, and 7; the three input points determining the first output point from the original second group of ten input points are now at locations 2, 5, and 8; and the three input points determining the first output point from the original third group of ten input points are now at locations 3, 6, and 9; and so forth. This interleave implies that sets of three adjacent output points use all different input points and do not require simultaneous memory accesses.
4. Consider 8-way parallelism and iMX, add more dummy outputs if necessary. See Figure 42d which shows the output points partitioned into four groups of 8 for parallel computations.
5. Compute coefficients and order as grouped. iMX will process one group at a time, using coefficient order from left-to-right, then up-to-down, then next group. Coefficients need to be arranged to the same order. If the iMX coefficient memory and the flash memory can accommodate all these coefficients, these coefficients can be included in the DSP code as constant data, and this step is done once in the software development. If the iMX coefficient memory can hold these coefficients all the time, but these take up too much room in the flash memory, this step can be performed once during system initialization. Likely the SDRAM can hold all these coefficients, but iMX coefficient memory cannot hold them all the time. this step should be performed once in the system initialization, and the coefficient image should be stored in SDRAM. When needed, these coefficients are swapped in from the SDRAM. If it is not desirable to store all these coefficients at any time, especially when M is very large (100+), compute needed "window" of coefficients with DSP concurrently with iMX processing. Just make sure the iMX coefficient memory can hold the necessary coefficients for a computation block.
6. Start computation on iMX. In this case, it takes about 12 cycles in the inner loop to produce the 27 valid output points. Each iMX command can produce a 2-D output block, so producing 16×27 output points will take about $10 + 16^*12 = 202$ cycles.
7. When iMX is done, have DSP pick the correct output points. In this example, 276 points are picked out of every group of 32 output points. This task will be easier to code if the width of output matches or is a multiple of 3^*M . DSP only has to touch each valid output once, so the loading of the DSP should not be significant.

[0133] In vertical resizing, iMX works in SIMD mode. Every group of 8 adjacent data input are processed in parallel. Coefficient are used one value per cycle, and this value should apply to all color components. Even if resizing factors are the same for horizontal and vertical, how iMX uses coefficients is different, so there needs to be a separate vertical resizing coefficient storage (which takes $1/3$ of horizontal coefficients). See Figure 42e. Again, there is the option to keep all vertical coefficients in iMX, swap in and out, or have DSP compute on the fly. DSP may need to pick valid

output rows after iMX completes processing.

Tone-scaling preferred embodiments

5 [0134] Tone-scaling operates on the dynamic range of the luminance signal (or the color signals) of an image to make details more clear. For example, a picture taken against the light or in a very bright environment typically has high brightness levels. Tone-scaling commonly relies on luminance (or color) histogram equalization as illustrated in block form by Figure 43. Indeed, converter block 430 converts the input luminance levels (in the range 0 to 255 for 8-bit or 0 to 4095 for 12-bit) to output luminance levels in the same range using a look-up table. The look-up table
10 consists of the pairs that are the input level and the corresponding output level with the output levels calculated in histogram equalization block 432 as follows. First, find the cumulative distribution function of the input luminance levels of the image to which the tone-scaling will apply; that is, find $F(r)$ such that $F(r) = (\text{the number of pixels with level } \leq r) / (\text{total number of pixels in the image})$. Next, create the look-up table function $T(r)$ through multiplication of $F(r)$ by the maximum pixel level and round-off to the nearest integer. Then the look-up table is just the pairs of levels (r,s) where
15 $s = T(r)$. Figure 45 illustrates $T(r)$ for an under-developed image (the majority of pixels have a low level as reflected by the large slope of $T(r)$ for small r) in which fine details in dark parts are difficult to perceive. Also as Figure 45 shows for this under-developed image, the tone-scaling converts the level $r = 500$ to $s = 2000$; and thus in the tone-scaled image the differences of the luminance levels will be emphasized for the low levels and de-emphasized for the high levels. Thus the tone-scaling enhances detail in dark portions.

20 [0135] However, the tone-scaled image may look unnatural in that the colors are too clear, as if the tone-scaled image were painted in oil paints. Thus this tone-scaling is sometimes too strong for consumer use because of the unnatural character even if the fine details are clearer; although other applications such as medical and night vision demand the fine detail despite unnaturalness.

25 [0136] The preferred embodiments provide tone-scaling by using a linear combination of the histogram equalization function $T(r)$ and the original image level r . That is, for a parameter α with $0 \leq \alpha \leq 1$ define a tone-scaling function by

$$s = \text{Round}(\alpha T(r) + (1-\alpha)r)$$

30 where $T(r)$ is as previously described except that the round off to the nearest integer is not needed in the definition of $T(r)$ because of the subsequent multiplication by α plus addition of $(1-\alpha)r$ and round off. Figure 45 illustrates the preferred embodiment for $\alpha = 0.3$ between the curve $s = T(r)$ and the identity line $s = r$.

[0137] Figure 44 shows preferred embodiment tone-scaling in functional block form: again define a histogram equalization function $T()$ for the luminance (or color) levels in block 442, and then define the rounded-off linear combination
35 with weight α of $T()$ and the identity in block 444 to yield the final look-up table for the tone-scaling in converter 440. When the weight α equals 0, then there is no tone-scaling and a natural look, but when the weight α equals 1, the tone-scaling is with $T()$ and fine details are enhanced. The value of weight α can be selected according to the application. All of the computations are programmable.

40 Implementation details

[0138] Preferred embodiment hardware structures supporting the foregoing functions include the following.

SDRAM Controller

45 [0139] SDRAM controller block 110 acts as the main interface between SDRAM 160 and all the function blocks such as processors' (ARM 130, DSP 122), CCD controller 102, TV encoder 106, preview engine 104, etc. It supports up to 80 MHz SDRAM timing. It also provides low overhead for continuous data accesses. It also has the ability to prioritize the access units to support the real-time data stream of CCD data in and TV display data out. It also provides power
50 down control for external SDRAM. DSP 122 can inhibit CKE signal of SDRAM 160 during no data access.

[0140] SDRAM controller block 110 supports 16/64/128/256 MB SDRAMs, 32-bit width or 2 x 16-bit width SDRAMs, maximum 80 MHz (e.g., 10-80 MHz) operation, availability of word, half-word, or byte access (ARM), commands: mode setting, power down and self refresh, programmable refresh interval, 2 or 3 CAS latency can be selectable, 2 Chip
55 Select Output (maximum SDRAM size is 1G bit), authorizes and manages DMA transfers, manages the data flow between processors SDRAM, CCD data buffer to SDRAM, preview engine to SDRAM, burst compression to/from SDRAM, video encoder from SDRAM, OSD from SDRAM, ARM to/from SDRAM, DSP image buffer to/from SDRAM. Figure 12a shows the data flow managed by the SDRAM controller. The signals and priorities are:

Signal Name	Signal Description
Clk	SDRAM clock (10-80 MHz)
Req	Data read/write request signal
req_en	Request enable (acknowledge) signal from SDRAM Controller When the peripheral modules require a data IN/OUT, the req signal shall be asserted and when the req_en signal is asserted, the req signal shall be negated
Address	Start address of read or write CCDC, PREVIEW, BURSTC, ENC, OSD, DSP: 22-bit width ARM: 25-bit width
Odata	output data to SDRAM (32-bit)
Idata	Input data from SDRAM (32-bit)
Rw	Read or Write signal 0: Write / 1: Read
Dten	Data write enable signal for DSP IF
Ds	Bus Select (4-bit) for ARM IF

The Priority list of access units is as follows,

Priority	Access Unit
1(highest)	ENC out
2	CCD in
3	OSD out
4	PRVW in
5	BURST in
6	DSP I/O
7	ARM I/O

Preview engine

[0141] Figure 14 is a block diagram of preferred embodiment preview engine 104 which provides image data with YCbCr in 4:2:2 format from CCD raw data from CCD-controller 102 and has the following main functions.

- Available for both RGB CCDs and complementary (YeCyMgG) CCDs (Figures 7a-7b show these CCD patterns)
- Digital gain adjustment
- White balance
- Vertical and horizontal noise filter
- RGB gain adjustment for complementary CCDs
- Independent gamma correction for RGB colors
- YCbCr-4:2:2 formatted data output

[0142] Sync module 1402 generates control signals for other modules such as a sync signal for a starting point of an image and an enable signal for down sampling. In this module, no image processing is executed. White balance module 1404 executes digital gain adjustment and white balance for CCD raw data. CFA interpolation module 1406 has many important sub-modules such as a horizontal noise filter, a horizontal interpolation, a vertical noise filter, a vertical interpolation, a down sampling, etc. This module outputs RGB formatted data irrespective of CCD mode (RGB CCD or complementary CCD). RGB gain modules 1408 for complementary CCD allow adjustment to white balance by RGB color format for complementary CCD. Gamma correction modules 1410 execute gamma correction with an approximated gamma curve having 4 linear segments. This module exists for each color to permit the independent adjustment to RGB. RGB2YCbCr conversion module 1412 converts RGB formatted data into YCbCr formatted data and adjusts offsets of Cb and Cr. 4:2:2 conversion module 1414 converts YCbCr-4:4:4 formatted data into 4:2:2 format

and outputs them on a 32-bit data bus. SDRAM interface module 1416 communicates with SDRAM controller 110 (Figure 1b) and requests it to store YCbCr-4:2:2 formatted image data.

The following describes the modules.

[0143] White balance module 1404 executes digital gain adjustment and white balance for CCD raw data. Digital gain adjusts for total brightness of the image and white balance adjusts the ratio of colors existing in a CFA pattern.

[0144] Figure 8 is a block diagram of white balance module 1404. There are two multipliers for the two gain adjustments and clip circuits to reduce the size of circuits. A gain value for digital gain named PVGAIN in this figure uses data in a PVGAIN register, and white balance is selected automatically by setting the CFA pattern register.

[0145] CFA interpolation module 1406 includes both sub-modules for horizontal and vertical interpolation and for horizontal and vertical noise filtering, down sampling, color adjustment and complementary color to RGB color conversion. Figure 10a is a block diagram of CFA interpolation module 1406. Horizontal noise filter sub-module 1002 executes a three-tap low pass filter horizontal filter; see Figure 10b.

[0146] Horizontal interpolation filter sub-module 1004 prepares two types of filters and interpolates horizontally using one of them. The outputs signal "L" and "R" means a left data and a right data on the line. For example, a processed line starts the following CFA pattern, GBGBGBGBGB ..., the output signal "L" is G and "R" is B. Therefore, these two outputs change the colors each line. Horizontal down-sampling sub-module 1006 outputs only data on valid pixels based on register settings of horizontal decimation pattern. Vertical interpolation sub-module 1008 processes a three-tap vertical interpolation filter using two line-memories 1010 outside the preview engine module and outputs data of all colors existing in the CFA pattern. And this sub-module also executes a vertical noise filter. Color selection sub-module 1012 extracts data by each color in the CFA pattern and outputs RGB color formatted data in RGB CCD mode or complementary color formatted data in complementary CCD mode. In this figure, "g" signal is temporal data regarding G and used for recalculating R and B in the next color adjustment sub-module 1014. The color formatted data is processed color adjustment in color adjustment sub-module 1014 and the processing is different depending on CCD mode. This image processing from vertical interpolation sub-module 1008 to color adjustment sub-module 1014 has a strong correlation depending on CCD mode and vertical interpolation mode. Therefore, the processing should be considered as a sequence of vertical interpolation processing as described below. Comp2RGB conversion sub-module 1016 converts complementary color format into RGB color format in complementary CCD mode. In RGB CCD mode, the data bypass this sub-module.

[0147] The following sections describe these sub-modules.

[0148] Horizontal noise filter 1002 executes three-tap horizontal low pass filter and can reduce random noise effectively. Actually, when the center of data is set to X_0 , the following calculation is executed depending on the CFA pattern and its processed line.

$$\{ (X_{-2} + 2X_0 + X_2) / 4 \quad (\text{two colors in processed line})$$

$$X_0 = \{$$

$$\{ (X_{-1} + 2X_0 + X_1) / 4 \quad (\text{one color in processed line})$$

An on/off switching of this filter can be controlled by a register setting.

[0149] Figure 10b is a block diagram of horizontal noise filter sub-module 1002. The two types of filter are implemented by using two adders and a switch named "three_taps_sw" in this figure. If there is one color in the processed line, the switch is set to on (High in the figure). This switch is automatically controlled depending on a register setting of the CFA pattern and a position of the line in the processed image. Before the output, noise-filtered data or bypassed data is selected by a register setting.

[0150] In horizontal interpolation sub-module 1004, there are two modes of filtering and the data from horizontal noise filter 1002 is interpolated horizontally by either a two-tap or five-tap interpolation filter. The two-tap filter utilizes the average the two data at the adjacent pixels on the left and right to interpolate the center data. This mode is called "simple mode". The five-tap horizontal interpolation filter utilizes the information of another color on the processed line so that a false color around an edge in processed image can be reduced effectively. This mode is called "normal mode". These modes are selectable by a register setting. Actually, when the center of data is set to X_0 , the following calculation is executed depending upon the interpolation mode.

$$\{ (-X_{-2} + 2X_{-1} + 2X_0 + 2X_1 - X_2) / 4 \quad (\text{normal mode})$$

5

$$x_0 = \begin{cases} \\ (X_{-1} + X_1) / 2 \end{cases} \quad (\text{simple mode})$$

10

Figure 10c shows an example of this horizontal interpolation processing in RGB Bayer CCD mode. In this figure, interpolated data is represented by small letters.

[0151] Figure 10d is a block diagram of horizontal interpolation module 1004. Two adders, one subtracter and a filter mode switch are implemented for executing one of these two types of filters. The filter mode switch is controlled by setting a register.

[0152] Vertical interpolation sub-module 1008 processes either a two-tap or three-tap vertical interpolation filter using two line-memories outside the preview engine module and outputs the information of all colors existing in the CFA pattern. And this sub-module also executes a vertical noise filter. An image processing in this module is a little complicated and the outputs from this sub-module is varied depending on a processed line, CCD mode, CFA pattern, filter mode and noise filter on/off. As explained in the following, the image processing from vertical interpolation sub-module 1008 to color adjustment sub-module 1014 has a strong correlation and this processing flow of them should be considered as a sequence of vertical interpolation processing. Therefore, this sequence of the vertical interpolation processing is explained first. The sequence may be called "vertical interpolation sequence".

[0153] As with horizontal interpolation, vertical interpolation processing also has two types of interpolation mode, that is "simple mode" and "normal mode". An interpolation filter in simple mode utilizes the average two data at the next pixels on the upper and lower to interpolate the center of data. In normal mode, the processing differs between RGB CCD mode and complementary CCD mode. The interpolation filter in normal mode in RGB CCD mode utilizes the data of one of the others color same as horizontal interpolation filter. Actually, when the data of a certain color to be interpolated is set to X (mainly R,B) and the data of a color utilized as a reference is set to Y (mainly G), the following calculation is executed depending on the interpolation mode through this vertical interpolation sequence and it is the output from color adjustment sub-module:

$$x_0 = \begin{cases} (X_{-1} - Y_{-1} + X_1 - Y_1) / 2 + Y_0 & (\text{normal mode}) \\ (X_{-1} + X_1) / 2 & (\text{simple mode}) \end{cases}$$

Figure 10e shows an example of this vertical interpolation sequence for the RGB Bayer CCD pattern.

[0154] In complementary CCD mode, normal mode means "simple interpolation with color adjustment". That is, data of all colors which is processed by simple vertical interpolation is adjusted based on the formula in complementary color space. Actually, when the data of a certain color to be interpolated is set to X and the data of the others color is set to W, Y, and Z the following calculations are executed in normal mode in complementary CCD mode.

$$x_0 = \begin{cases} (X_{-1} + X_1) / 2 \pm a(w_0, x_0, y_0, z_0) & (\text{normal mode}) \\ (X_{-1} + X_1) / 2 & (\text{simple mode}) \end{cases}$$

As to the calculation of $a = a(w_0, x_0, y_0, z_0)$, see below.

[0155] In this vertical interpolation sequence, main roles of vertical interpolation sub-module 1008 are to execute a part of vertical interpolation sequence and vertical noise filter. The part of vertical interpolation sequence means preparing data for normal vertical interpolation mode. As shown in Figures 10e and 10b (for RGB and complementary

CCD patterns, respectively), in simple mode an output data of this vertical interpolation sub-module bypasses color adjustment sub-module. Therefore, in simple mode, the output from this sub-module is used as the output of vertical interpolation sequence. In any case of interpolation mode, this sub-module calculates the following equation for vertical interpolation sequence.

$$x_0 = (X_{-1} - X_1) / 2$$

Vertical noise filter ... which executes the following 3 taps vertical low pass filter is also processed in this sub-module depending on the CFA pattern.

$$x_0 = (X_{-1} - 2X_0 - X_1) / 4$$

However, for this filtering, data of same color on processed 3 lines must be prepared. Therefore, a function of the vertical noise filter mainly executes only G in RGB Bayer CCD. Figure 10g shows an example of the output of this vertical interpolation sub-module for a RGB Bayer CCD. When the vertical noise filter can be applied and it is set on, original data (R in this figure) is also adjusted in order to keep a correlation to the others color (G in this figure).

[0156] Figure 10h is a block diagram of vertical interpolation sub-module 1008. Six adders and two subtractors are implemented for executing vertical interpolation and noise filtering. Especially, a calculation process of L_121 and R_121 is so complicated that switching operation for L_121 and R_121 is not shown to simplify this figure.

[0157] Color selection sub-module 1012 arranges the inputs from vertical interpolation sub-modules in order of color format, that is R, G and B in RGB CCD mode or Ye, Cy, Mg, G in complementary CCD mode. This arrangement is executed automatically by setting register of the CFA pattern. Figure 10i shows an example of this color selection processing in RGB Bayer CCD of Figure 10g. The outputs named "g" in this figure is a temporal data of G and is used for recalculation of R or B in RGB CCD mode in color adjustment sub-module.

[0158] Figure 10j is a block diagram of color selection sub-module 1012. Four color extractors switch and select independently correct colors from four inputs from vertical interpolation sub-module 1008.

[0159] Color adjustment sub-module 1014 executes the rest of calculation for vertical interpolation sequence. In RGB CCD mode such as RGB Bayer CCD, R or B is recalculated using the temporal data of G. When data of R or B from color selection sub-module is set to X, the following calculation is executed in RGB CCD mode.

$$x = X - G_{temp} + G$$

In the example of Figure 10i, when noise filter is off,

$$X = (b_{02} - b_{22}) / 2$$

$$G_{temp} = (G_{02} + G_{22}) / 2$$

$$G = g_{12}$$

Therefore,

$$\begin{aligned} x &= B \\ &= (b_{02} - b_{22}) / 2 - (G_{02} + G_{22}) / 2 + g_{12} \\ &= ((b_{02} - G_{02}) + (b_{22} - G_{22})) / 2 + g_{12} \end{aligned}$$

This is the output B of the color adjustment module and also the output of vertical interpolation sequence. That is, vertical interpolation sequence in RGB CCD mode utilizes the average of differences between data of color to be interpolated and reference data of the others color.

[0160] In complementary CCD mode, color adjustment is processed to data of all colors from color selection sub-module. First, value a is calculated at each pixel based on a formula in complementary color space $Ye + Cy = G + Mg$.

5

$$a = G + Mg - Ye - Cy$$

That is, the value a can be considered as the amount of an error value of four colors. Therefore, in complementary CCD mode, to data of all colors, Ye , Cy , Mg and G , the following adjustment is processed to satisfy the above formula.

10

$$ye = Ye + a/4$$

$$cy = Cy + a/4$$

15

$$g = G - a/4$$

20

$$mg = Mg - a/4$$

Figure 10k is a block diagram of color adjustment sub-module 1014. Six adders and three subtractors are implemented for executing the two types of calculations described above. A switcher named CCDMOD in this figure selects correct outputs depending on CCD mode and is controlled by setting a register.

25

[0161] Comp2RGB conversion sub-modules 1016 converts complementary color formatted data to RGB formatted data in complementary CCD mode. Especially for G , data from color adjustment and data calculated by conversion formula can be blended by 5 types of blending ratio. Actually, the following calculation is executed based on the conversion formula:

30

$$R = Ye - Cy + Mg$$

$$G = rG_{\text{input}} + (1-r)(Ye + Cy - Mg) \quad (r=0, 1/4, 2/4, 3/4, 1)$$

35

$$B = Mg - Ye + Cy$$

In RGB CCD mode, data from color adjustment sub-module bypass this sub-module.

40

[0162] Figure 101 is a block diagram of comp2RGB conversion submodule 1016. Three adders, three subtractors, and two multipliers are implemented for executing the calculations above. A gain adjuster for G named "green_ratio" in this figure is adjustable by setting a register. In RGB CCD mode, a CCDMOD switcher selects off (high in this figure) for bypassing this module.

45

[0163] RGB gain for complementary CCD module allows adjustment of white balance by RGB color format even for complementary CCD module. This module is also available in RGB CCD mode.

[0164] Figure 9a is a block diagram of complementary white balance module 1408. One multiplier and clip circuit is implemented for this operation. Each gain for RGB is set by a register.

50

[0165] Gamma correction modules 1410 execute gamma correction for each color data in RGB color format. For this operation, prepare in advance three types of data for approximating the gamma curve by four linear segments. Those are area, offset and gain shown in Figure 9b. As shown in Figure 14, this module exists for each color so that the independent adjustment to RGB may be made.

[0166] Figure 9c is a block diagram of gamma correction module 1410. Area detector selects correct gain and offset for input data based on area data. The data regarding gain, offset, and area are set in three registers.

55

[0167] RGB2YCbCr conversion module 1412 converts RGB formatted data to YCbCr formatted data and adjusts offsets to Cb and Cr based on the following matrix calculation.

$$\begin{array}{c}
 \left[\begin{array}{c} Y \\ Cb \\ Cr \end{array} \right] \left[\begin{array}{ccc} COEF1 & COEF2 & COEF3 \\ COEF4 & COEF5 & COEF6 \\ COEF7 & COEF8 & COEF9 \end{array} \right] \left[\begin{array}{c} R \\ G \\ B \end{array} \right] + \left[\begin{array}{c} 0 \\ OFFSET_Cb \\ OFFSET_Cr \end{array} \right]
 \end{array}$$

Each coefficient in this matrix is set by a register so that variable setting for this conversion is available.

[0168] Figure 11 is a block diagram of this RGB2YCbCr conversion module 1412. Nine multipliers and five adders are implemented for the foregoing matrix calculation. After multiplying RGB data with coefficients, the six least significant bits of each data from the multipliers is cut in order to reduce size of circuits. As to Cb and Cr, additional circuit for offset adjustment is followed by YCbCr conversion circuit. Clip circuits for Cb and Cr includes conversion circuits from two's complement to offset binary.

[0169] Burst mode compression/decompression engine

[0170] The preferred embodiment DSC engine includes an improved Burst Capture function with real-time processing, without compromise in the image resolution as compared to the regular capture mode. The Burst Capture Mode is the use of dedicated compression and decompression engine 108 for an increased burst capture sequence length. A sequence of CCD raw image frames is first stored in SDRAM 160 by using Compression engine 108. Then, as an off-line process, the image pipeline of regular capture mode retrieves the CCD raw images from SDRAM 160, processes them sequentially, and finally stores them back as JPEG files in the SDRAM. The Animated Playback Mode can display these JPEG files.

[0171] Burst mode compression/decompression engine 108 includes differential pulse code modulation (DPCM) and Huffman coding using the same tables as the entropy-coding of DC coefficients in baseline JPEG compression. Engine 108 uses the sample Huffman table in the JPEG standard for chrominance DC differential data. Engine 108 also provides the inverse transforms as illustrated in Figure 13. Fixed Huffman Table (JPEG Huffman table for Chrominance DC coefficients):

Categor Y (SSSS)	\hat{D}_i	Code Length	Codeword
0	0	2	00
1	-1,1	2	01
2	-3,-2,2,3	2	10
3	-7,...,-4,4,...,7	3	110
4	-15,...,-8,8,...,15	4	1110
5	-31,...,-16,16,...,31	5	11110
6	-63,...,-32,32,...,63	6	111110
7	-127,...,-64,64,...,127	7	1111110
8	-255,...,-128,128,...,128	8	11111110
9	-511,...,-256,256,...,511	9	111111110
10	-1023,...,-512,512,...,1023	10	1111111110
11	-2047,...,-1024,1024,...,2047	11	11111111110
12	-4095,...,-2048,2048,...,4095	12	111111111110

The encoder has four look-up tables: Huffman code (13 x 2-byte entries), Huffman code length table (13 x 1-byte entries), low bit mask to generate variable-length bit stream (32 x 4-byte entries), and log table (256 x 1-byte entries). The Huffman tables are not programmable for simplicity, although alternative embodiments could include programmable Huffman tables.

[0172] The Huffman decoder performs the inverse function of the Huffman encoder and has five look-up tables: max code comparison table (13 x 2-byte entries), Min code comparison table (13 x 2-byte entries), decoded Huffman symbol pointer (13 x 1-byte entries), decoded Huffman symbol table (13 x 1-byte entries), and bit position mask (32 x 4-byte entries).

[0173] The lossy mode compression just discards the least significant bit (LSB) or the two least significant bits of

each coefficient.

Playback synchronization

5 [0174] A problem involved in playback of audio-visual bitstreams is how to synchronize audio with video signal. The preferred embodiments play the audio bitstream seamlessly in the background in real-time with the audio encoded by using the simple coding standards like ITU-T G.711 and Microsoft 16-bit PCM. By using an interrupt service routine, about 0.1% of the DSP resources is enough to output audio in real time through (multichannel) buffered serial ports; see Figure 1b. Therefore, the preferred embodiment must realize the video decoding in synchronization to the audio playback.

10 [0175] For clarity, assume that both audio and video are captured in full speed (real-time with 8K sample/s for audio and 30 frame/s for video). Audio is played back as samples. However, video is displayed in the granularity of frames. Thus the synchronization problem is caused by the fact that the video decoding could be faster or slower than the real-time requirement. If the video decoding is too fast, a certain amount of delay slots has to be inserted to slow down the decoding. Contrarily, if the video decoding is too slow, some video frames must be skipped to catch up with the real-time audio playback.

15 [0176] The preferred embodiments handle both cases. Especially in the case of slow video decoding, the preferred embodiments can properly select and skip the frames in an optimal manner. Note that the preferred embodiment is described for video bitstreams without bi-directional coded frames (B-frames).

20 [0177] Figure 46a depicts the synchronization between audio and video. The first video frame is pre-decoded before beginning audio-video playback. Since the video is displayed in the granularity of frames, the synchronization points are located at the video frame boundaries, i.e. $\{t = 0, \Delta T, 2\Delta T, 3\Delta T, \dots\}$. Here ΔT is the duration of a frame, which is defined as:

$$25 \quad \Delta T = 1/f_p \quad (1)$$

where f_p is the frame-rate used for the video sequence.

30 [0178] Audio and video could lose synchronization when the video decoding speed is not fast enough. As illustrated in Figure 46a, when the decoding of video frame 2 has not finished in time ($T_{d2} > \Delta T$), the audio-video playback loses synchronization after displaying video frame 1. Here $\{T_{dm}, m=0, 1, 2, \dots\}$ denotes the decoding time used for decoding video frame m .

35 [0179] With insufficient video playback speed, the only way to maintain a reasonable synchronization between audio and video is to skip video frames properly. In Figure 46b, video frame 2 is skipped (and frame 1 repeated) so that synchronization can be reacquired at frame 3.

40 [0180] A preferred embodiment circular buffer scheme is illustrated in Figure 47. The video decoder is connected to one side of the circular buffer, the display is connected to the other side. The circular buffer has a size of N video frames. There are two registers associated with each frame buffer of the circular buffer: the first register contains TP_n , $n=0, 1, 2, 3, \dots, N-1$ which indicates the presumptive presentation time of the video frame stored in buffer n , and the second register contains S_n , $n=0, 1, 2, 3, \dots, N-1$ which signals whether the frame in buffer n is ready for display (1 for ready, 0 for not ready). Of course, the value of TP_n is a multiple of ΔT . Buffer switching for display also occurs at frame boundaries (i.e. at time $t = m\Delta T$, $m=0, 1, 2, \dots$). Because the preferred embodiments use a circular buffer containing N frames, all the indices $(\dots, n-1, n, n+1, \dots)$ should be regarded as modulo- N indices.

45 [0181] Suppose the time after decoding the current video frame is T . The decoded current frame is stored in buffer $n-1$ in Figure 47. Therefore, the buffer to be used for storing the next frame in Figure 47 is buffer n .

[0182] Determine the current position in the bitstream: the frame index m of the current decoded frame is defined as

$$50 \quad m = TP_{n-1} / \Delta T \quad (2)$$

Determine the decoding starting time of the next frame: since the frame in the buffer n is to be displayed during the time interval of $\{TP_n \leq t < TP_{n+1}\}$, buffer n is not available for decoding the next frame until TP_{n+1} . Therefore, the decoding starting time of the next frame T_s is:

$$55 \quad T_s = \max\{T, TP_{n+1}\} \quad (3)$$

Determine the next frame to be decoded: let \hat{Td} be the estimated time for decoding the next frame, the presentation time of the next frame must satisfy:

5

$$\begin{cases} TP_n > Ts + \hat{Td} \\ TP_n \geq TP_{n-1} + \Delta T \end{cases}$$

10

The above conditions imply that the decoding of the next frame is finished before its presentation time, and the next frame is located at least a frame after the current frame in the bitstream.

Because TP_n must be a multiple of ΔT , the next frame that can be synchronized to audio satisfies the conditions:

15

$$\begin{cases} TP_n = \Delta T \left[\frac{Ts + \hat{Td}}{\Delta T} + 0.5 \right] \\ TP_n \geq TP_{n-1} + \Delta T \end{cases}$$

20

where $[\cdot]$ denotes integer part by truncation.

[0183] Therefore, the presentation time of the next frame is determined by:

25

$$TP_n = \max \left\{ \Delta T \left[\frac{Ts + \hat{Td}}{\Delta T} + 0.5 \right], TP_{n-1} + \Delta T \right\} \quad (4)$$

30

There are different methods to estimate \hat{Td} , such as using statistical estimation based on prior decodings or frame parameters. One preferred embodiment simply uses the actual decoding time of the most recently decoded frame of the same picture coding type (I-frame or P-frame) plus a certain amount of safety margin as the estimated decoding time for the next frame.

[0184] The frame index m' of the next frame to be decoded can thus be computed as:

35

$$m' = TP_n / \Delta T \quad (5)$$

Then the number of frames Δm to be skipped from the current position is determined by:

40

$$\Delta m = m' - m - 1 \quad (6)$$

Equation (2) to (6) make up of the basic control operations for updating the circular buffer.

45

[0185] The preferred embodiments use the circular buffer scheme to realize synchronization. There are two parts: the video decoder buffer switch control and the display buffer switch control. Figure 48 demonstrates the flowchart of the video decoder buffer switch control, which contains two stages: initialization and playback.

[0186] Initialization: in the circular buffer initialization, N_i ($1 \leq N_i \leq N$) video frames are decoded before starting playback. As shown in the dashed box in Figure 48, there are four steps for the initialization:

50

step 0: set all the presentation time registers $\{TP_n, n=0, 1, 2, 3, \dots, N-1\}$ and the status registers $\{S_n, n=0, 1, 2, 3, \dots, N-1\}$ to zero, switch the video decoder to buffer 0 (i.e. $n=0$), point to the beginning of the video bitstream (i.e. $m' = \Delta m = 0$), set time to zero (i.e. $t=0$)

55

step 1: set the related status register S_n to 1, skip Δm video frames, decode frame m' , store the decoded frame in buffer n . (Recall on the first pass through the loop, $n=0, m'=0$, so the first frame is decoded and stored in buffer 0.)

step 2: set the decoding start time Ts to t , switch to the next buffer (i.e. $n++$), update $TP_n, m', \Delta m$ according to equations (4), (5), and (6).

step 3: check whether the number of decoded frames reaches the pre-set frame number N_f . If true, go to playback, otherwise, loop to step 1.

[0187] Playback: there are six steps involved in updating the circular buffer during the playback.

5

step 0: switch display to buffer 0, enable display, reset time to zero (i.e. $t = T = 0$), switch the video decoder to buffer N_f (i.e. $n = N_f$)

step 1: if the whole video sequence is decoded, stop decoding, otherwise, go to step 2.

step 2: update T_s , TP_n , m' and Δm according to equations (3), (4), (5), and (6).

10

step 3: wait until time reaches T_s (i.e. $t \geq T_s$), go to step 4.

step 4: set the related status register S_n to 0, skip Δm video frames, decode frame m' , store the decoded frame in buffer n .

15

step 5: if the frame decoding finishes in time (i.e. $t < TP_n$), set S_n to 1 to indicate the decoded frame is ready for display, set T to t , switch the video decoder to the next buffer (i.e. $n++$). Otherwise, set T to t , add DT to the estimated \hat{T}_d (i.e. $\hat{T}_d = DT$ with $DT = N_d \Delta T$, intentionally skip N_d ($0 \leq N_d$) more frames in the next stage). set the current frame index m to m' . Go to step 1. Note that N_d is a parameter to control the screen freezing time before resuming the synchronization.

20

[0188] Users can freely decide the circular buffer size (N), the initial time delay (N_f) for the playback as well as the screen freezing time (N_d). Obviously, the minimum buffer size is 3 video frames (i.e. $N = 3$), the least time delay is one video frame (i.e. $N_f = 1$). However, in the case of insufficient video decoding speed, it is strongly recommended to decode $N-1$ frames (i.e. $N_f = N-1$) during the circular buffer initialization, so that the video decoder can gain the maximal room to catch up with the audio real time playback.

25

[0189] Display buffer switch control: the display buffer switch control is carried out in parallel to the video decoder buffer switch. The preferred embodiment checks the display buffer switch at video frame boundaries: $t = m\Delta T$, $m=0, 1, 2, \dots$. Suppose the display is currently showing the video frame in buffer $n-1$, it switches to the next buffer, i.e. buffer n , if and only if the current time ($t \geq TP_n$) and ($S_n=1$) holds. Otherwise, it is connected to buffer $n-1$. Here, if ($t \geq TP_n$) and ($S_n=0$), it means the decoder has not finished decoding of the frame in time. In this case, the video frame in buffer n has been discarded, the decoder is decoding the conservatively selected next frame to update buffer n again, the display should keep displaying the frame in buffer $n-1$ until ($t \geq TP_n$) and ($S_n=1$) holds.

30

[0190] In summary, the preferred embodiment provides a way to realize the synchronization between audio and video when playing back by using software or firmware.

Variable length decoding

35

[0191] Variable Length Decoding (VLD) is involved in decoding bitstreams which are generated by using Variable Length Encoding (VLC) at encoder; see Figure 1b item 126. Because of VLC, the number of bits used for coding units varies from unit to unit. Therefore, a decoder does not know the number of bits used for a coding unit before having decoded it. This makes it essential for a decoder to use a bitstream buffer during the decoding process.

40

[0192] In video coding, for example, a frame to be encoded is decomposed into a set of macroblocks (see Figure 49). Under the consideration of the smallest memory requirement, a coding unit here is normally defined as macroblock, which consists of a 16×16 pixel luminance area and the corresponding chrominance areas depending on the chroma format (4:2:0, 4:2:2, or 4:4:4). Certainly, a slice (a row of macroblocks in a frame) or even the frame itself can be treated a coding unit if there is enough memory.

45

[0193] Figure 50 depicts the video playback on a preferred embodiment digital still camera (DSC). In DSC applications, the video bitstream is pre-captured and stored on the high-capacity SDRAM, and the video decoder is built on the DSP. Since it is extremely expensive for the decoder to directly access the SDRAM, an on-chip bitstream buffer is opened on the DSP internal memory. The bitstream is first loaded from SDRAM to the bitstream buffer through the SDRAM, then the decoder uses the bitstream in the bitstream buffer to reconstruct video. Since the bitstream loading is achieved by using DMA (Direct Memory Access), which can run in the background without intervention of a CPU, the bitstream loading overhead is mainly due to time used for setting up registers for the DMA transfer.

50

[0194] There are two basic requirements in terms of bitstream buffer management. First of all, the buffer size should be big enough to cover the worst case. For example, in video coding, the theoretically maximal number of bits for encoding a macroblock could be 256 words (one word here is defined as two bytes). Although this worst case is very rare, the bitstream buffer size has to be 256 words in order to be at the safe side. Secondly, the bitstream buffer should never underflow, that is, the buffer management should guarantee that the bitstream for a coding unit is available when it is being decoded.

55

[0195] There are different schemes to satisfy the second requirement. The simplest one would be to check the

decoding position in the bitstream buffer at each buffer access. The bitstream buffer is re-filled whenever the decoding position is out of the valid buffer range. Because the decoding is a bit by bit operation, this scheme is not realistic: it spends too much overhead in deciding when to re-fill the buffer.

[0196] A realistic scheme is the linear shifting buffer scheme as shown in Figure 51a. In this scheme, the bitstream buffer is linearly accessed by the decoder from left to right, after decoding a unit the rest of the bitstream is shifted forward to the beginning of the buffer, then the buffer is re-filled to "full" before decoding the next unit. In Figure 51a, Ps and Pd denote the current decoding position and the bitstream end position in the bitstream buffer, respectively.

[0197] This buffer scheme has two disadvantages. First, since the buffer size is much larger than the average number of bits of the decoding units, a lot of time will be spent on the bitstream shifting. For instance, in video decoding the buffer size is 256 words to cover the worst case, but on average a unit may only use 16 words, this means about 240 words of shifting for each unit. The second disadvantage is that it requires a bitstream loading after decoding each unit; this costs additional overhead because time has to be spent on issuing the DMA transfers.

[0198] A better buffer management scheme is so-called quasi-circular buffer scheme as shown in Figure 51b. In this scheme, the decoder accesses the bitstream buffer in a circular manner. This avoids the bitstream shifting required by the linear buffer scheme. There are two cases after decoding a unit. This first case is in the lefthand portion of Figure 51b: the rest of bitstream is located in the middle of the buffer. In this case, the buffer is filled by loading the bitstream twice, one for the right end followed by the other one for loading the left end. (Note: if the bitstream loading can write the bitstream into the bitstream buffer in a circular manner, only one load is needed; however, this is not always the case.) The second case is shown in the righthand portion of Figure 51b, in which only the middle of the buffer needs to be filled.

[0199] The quasi-circular buffer scheme is much more efficient than the linear shifting buffer because it avoids bitstream shifting, but it still suffers from a disadvantage that one or two bitstream loads are needed after decoding each unit. The following preferred embodiment hybrid circular-double buffer scheme solves this problem.

[0200] Figure 52 status 0 shows a hybrid circular-double buffer containing two buffers of equal size; namely, the left buffer and the right buffer. There is a flag for each buffer to indicate the buffer fullness ("full" / "not-full"). Ps points to the current decoding position after decoding a unit. In terms of buffer size, each buffer covers the worst case of decoding coding units, this makes the hybrid buffer size twice of a linear shifting buffer or a quasi-circular buffer. Unlike a traditional double buffer, the two buffers here have a continual memory allocation, i.e. the left buffer is directly followed by the right buffer in the memory map. The decoder accesses the hybrid buffer in a circular manner.

[0201] The preferred embodiment hybrid buffer operates through the following four statuses:

[0202] • Status 0: the initialization status, both the left and right buffers are fully loaded and set to "full", Ps points to the beginning of the hybrid buffer.

Status 1: after decoding the first unit, change the left buffer flag to "not-full".

Status 2: after decoding a unit, if the current decoding position Ps is in the right buffer and the left buffer flag is "not-full", fully load the left buffer and set the left buffer flag to "full". In addition, if the right buffer flag is "full", change it to "not-full". Otherwise, no action is taken.

Status 3: after decoding a unit, if the current decoding position Ps is in the left buffer and the right buffer flag is "not-full", fully load the right buffer and set the right buffer flag to "full". If the left buffer flag is "full", change it to "not-full". Otherwise, no action is taken.

[0203] Taking the preferred embodiment platform (e.g., Figure 1b) as an example (where data is in 16-bit units), define the following data type:

```
typedef struct bitstream {
    SInt bit_ptr; /* current bit position (0 ~ 16)
    */
    SInt Ps; /* current decoding position in
    bitstream buffer */
```



```

    Sint left_flag      /* left buffer flag "full / not-
5      full"           */
    Sint right_flag     /* right buffer flag "full / not-
10     full"           */
    USInt *databuf;     /* bitstream buffer
                        */
    Long Addr_SDRAM;    /* bitstream address in SDRAM
                        */
15     } Bitstream;

```

20 The pseudo code shown in Table 1. describes the hybrid circular-double buffer scheme. Function BufferInitialization() is called only once at the beginning of decoding, while function BitstreamBufferUpdate() is called after decoding each coding unit. it automatically updates the buffer flags and re-loads the buffers if the conditions become true. In Table 1 BUFSIZE stands for the buffer size of the hybrid circular-double buffer.

```

25     Void BufferInitialization(
        Bitstream *stream, /* pointer of bitstream
                           */
30     )
    [
35     /*=====
        =====*/
        /* Initialization of the hybrid circular-double buffer
40     */
        /*=====
        =====*/
45     LoadBuffer(&stream->databuf[0],      stream->Addr_SDRAM,
        BUFSIZE);
        stream->Addr_SDARM +=BUFSIZE;
50
55

```



```

    stream->left_flag = "full";
    stream->right_flag = "full";
5    stream->Ps = 0;
    stream->bit_ptr = 16;
    }
10
15
Void BitstreamBufferUpdate(
    Bitstream *stream, /* pointer of bitstream
                        */
20    )
{
25    /*=====
=====*/
    /* Update the left buffer if necessary
30    */
    /*=====
=====*/
35    if (stream->left_flag == "not-full"    && stream->Ps >=
BUFSIZE/2)
    {
40        LoadBuffer(&stream->databuf[0], stream->Addr_SDRAM,
BUFSIZE/2);
        stream->Addr_SDRAM +=BUFSIZE/2;
45        stream->left_flag = "full";
    }
50
    /*=====
=====*/
55

```



```

5      /* Update the right buffer if necessary
                                           */
      /*=====
=====*/
10      if (stream->right_flag == "not-full"    && stream->Ps <
BUFSIZE/2)
      {
15          LoadBuffer(&stream->databuf[BUFSIZE/2],      stream-
>Addr_SDRAM, BUFSIZE/2);
          stream->Addr_SDARM +=BUFSIZE/2;
          stream->right_flag = "full";
20      }

25      /*=====
=====*/
      /* Update the left buffer flag
                                           */
30      /*=====
=====*/
35      if (stream->left_flag == "full"    && stream->Ps <
BUFSIZE/2)
          stream->left_flag = "not-full";
40

45      /*=====
=====*/
      /* Update the right  buffer flag
                                           */
50      /*=====
=====*/
      if (stream->right_flag == "full"    && stream->Ps >=
55      BUFSIZE/2)

```


5

10

```

stream->right_flag = "not-full";
}

```

15

Table 1. Pseudo code for the hybrid circular-double buffer scheme

As it can be seen in BitstreamBufferUpdate() in Table 1, the left buffer or right buffer is not reloaded after decoding each unit, but is loaded only if the opposite buffer (left / right) is in use and its buffer flag is "not-full". This greatly reduces the number of buffer loads. Consider the video coding as an example. This needs BUFSIZE of 512 words if a macroblock is the unit, the average bitstream size of a unit is assumed to be 16 words. Because the linear shifting buffer and the quasi-circular buffer re-fill the buffer after decoding each unit, the average loading length for those two schemes is also 16 words. Compared with the fixed loading length of 256 words in the hybrid circular-double buffer scheme, the preferred embodiment reduces the loading overhead by a factor about 16 (i.e. 256 / 16).

[0204] Mini-experiments compared the three buffer schemes discussed above. The video sequence used was coast-guard (352 x 288, 300 frames, 4:2:0). The bitstream is generated by using a MPEG1 video encoder. The target bit-rate is 3 Mbit/s, I-frame only. The same decoder with three different buffer schemes are used to decode the same bitstream, the buffer loading count and word shifting count are recorded during the decoding. The performance comparison among the three buffer schemes is listed in Table 2. As shown in Table 2, for each macroblock the linear shifting buffer scheme requires one buffer load, and on average about 240 words of shifting. The quasi-circular buffer scheme needs slightly more buffer loads (1.06 load/macroblock) but no shifting. The preferred embodiment hybrid circular-double buffer scheme used only about 0.0619 buffer load per macroblock. On the preferred embodiment platform of Figure 1b in particular, the preferred embodiment scheme provides a cycle count reduction ratio of about 113 and 17 in comparison to the linear shifting buffer scheme and the quasi-circular buffer scheme, respectively.

35

Table 2.

Performance comparison among three buffer schemes on TMS320DSC21 platform			
	Linear shifting buffer	Qu-circular buffer	Hybrid circular-double buffer
Buffer size (words)	256	256	512
Number of loads per macroblock	1.00	1.06	0.0619
Number of word shifting per macroblock	240.15	0	0
Overhead per load (cycles)	80	80	80
Cycle count per word shifting	2	2	2
Total cycles used for bitstream buffer per macroblock	560.30	84.72	4.95
Cycle count ratio vs. the hybrid circular-double buffer scheme	113.19	17.12	1.00

50

Onscreen display and graphics acceleration

55

[0205] The Onscreen display (OSD) module 105 is responsible for managing OSD data from different OSD windows and blending it with the video. It reads OSD data from SDRAM 160, and outputs to NTSC/PAL encoder 106. The OSD module defaults to standby mode, in which it simply sends video to NTSC/PAL encoder 106. After being configured and activated by ARM CPU 130, the OSD module reads OSD data and mixes it with the video output. ARM CPU 130

is responsible for turning on and off OSD operations and writing the OSD data to the SDRAM. Figure 15 shows the block diagram of the OSD module and affiliated other items. The various functions of the OSD are described in the following paragraphs.

[0206] OSD data storage. The OSD data has variable size. In the bitmap window, each pixel can be 1, 2, 4, or 8 bits wide. In the YCrCb 4:2:2 window, it takes 8-bit per components, and the components are arranged according to 4:2:2 (Cb/Y/Cr/Y ...) format. In the case where RGB graphics data needs to be used as OSD, the application should perform software conversion to Y/Cr/Cb before storing it. The OSD data is always packed into 32-bit words and left justified. Starting from the upper left corner of the OSD window, all data will be packed into adjacent 32-bit words.

[0207] Setting up an OSD window. An OSD window is defined by its attributes. Besides storing OSD data for a window into SDRAM by ARM CPU 130, the application program also needs to update window attributes and other setup in the OSD module as described in the following subsections.

[0208] Location register. The Location register contains X and Y locations of the upper left and lower right corners of each window. The application program needs to set up the CAM and enable selected OSD windows; see Figure 16.

[0209] Color look up tables. The OSD has the fixed 256-entry color look up table (CLUT). The CLUT is used to convert bitmap data into Y/Cr/Cb components. In the case of 1.2 or 4 bitmap pixels, the CLUT can be determined by CLUT registers.

[0210] Blending and transparency. Color blending on the pixel level is also supported. This feature is available for the bitmap displays only (Window1,2). If the window color blending is enabled, the amount of blending of each pixel is determined by the blending factor. As shown in the following table, the window blending supports 5 different levels, according to the selected blending factor. The hardware also supports a transparency mode with bitmap. If transparency is enabled, then any pixel on the bitmap display that has a value of 0 will allow video to be displayed. Essentially, 0-valued pixels are considered the transparent color, i.e. the background color will show through the bitmap. The Table shows the connection between transparency and blending on the same window.

Transparency	Blend Factor	OSD window contribution	Video contribution
OFF	0	0	1
	1	1/4	3/4
	2	1/2	1/2
	3	3/4	1/4
	4	1	0
ON	0	if pixel value = 0	if pixel value = 0
	1	0	1
	2	1/4	3/4
	3	1/2	1/2
	4	3/4	1/4
		1	0

[0211] Hardware cursor. A rectangular shape is provided using hardware window1. With window1, the cursor always appears on top of other OSD Windows. The user can specify the size, color of the shape. When hardware window1 is designated as the cursor, only two windows are available for the OSD application. If a hardware cursor is not used, then the application can use window1 as a regular hardware window. Figure 17 shows an example of the hardware cursor.

DSP subsystem

[0212] The DSP subsystem consists of C54x DSP, local memory blocks, iMX and VLC accelerators, shared image buffers, and the multiplexers implementing the sharing.

[0213] C54x is a high performance, low power, and market proven DSP. cDSP hardware and software development tools for C54x are also very mature.

[0214] The DSP carries out auto exposure, auto focus, auto whitebalancing (AE/AF/AWB) and part of the image pipeline tasks. It also handles SDRAM transfer and drives the accelerators to implement the rest of image processing and image compression tasks. Flexibility and ease of programming in the DSP enables camera makers to refine the image processing flow, adjust quality/performance tradeoffs, and introduce additional features to the camera.

[0215] The configurable DSP (cDSP) design flow is adopted to allow flexibility and design reuse. The memory blocks time-shared among DSP and accelerators are large enough for one processing unit (16x16 pixels) and provide zero-wait state access to DSP.

Features

Fixed-point Digital Signal Processor

100 MIPs LEAD2.0 CPU

On-module RAM 32Kx16bit

5 (4 blocks of 8Kx16bit dual access program/data RAM)

Multi-Channel Buffered Serial Ports (McBSPs)

ARM can access RAM via Enhanced 8-bit Host Port Interface

One hardware timer

On-chip Programmable PLL

10 Software Programmable Wait-State Generator

Scan-based emulation and JTAG boundary scan logic

[0216] Figure 18a shows more details on the DSP subsystem and in particular the details of the connection between the DSP and the iMX and VLC. Figure 18b is the memory map.

15 [0217] The shared memory blocks A and B occupy two 2Kword banks on the DSP's data memory space. Each block can be accessed by DSP, iMX, VLC, and SDRAM controller depending on static switching controlled by DSP. No dynamic, cycle-by-cycle, memory arbitration is planned. DSP's program should get seamless access of these memory blocks through zero-wait-state external memory interface.

20 [0218] The configuration memory blocks, for iMX coefficient, iMX command, VLC Q-matrix, and VLC Huffman table, also connect to DSP's external memory interface. They are also statically switched between the specific module and DSP. Typically at powerup or at initial stage of camera operation mode, these memory blocks are switched to DSP side so DSP can set up the appropriate configuration information for the operation. Then, they are switched over to iMX and VLC for the duration of operation.

25 Imaging Extension (iMX)

[0219] iMX, imaging extension, is a parallel MAC engine with flexible control and memory interface for extending image processing performance of programmable DSPs. iMX is conceived to work well in a shared memory configuration with a DSP processor, such that flexibility, memory utilization, and ease of programming are achieved. The architecture covers generic 1-D and 2-D FIR filtering, array scaling/addition, matrix multiplications (for color space transform), clipping, and thresholding operations.

[0220] For digital still cameras, iMX can be used to speed up

35 CFA interpolation,
color space conversion,
chroma down-sampling,
edge enhancement,
color suppression,
DCT and IDCT,
40 Table lookup.

[0221] iMX methodology originates from the discipline of parallel processing and high performance computer architecture. The design comprehends the need for a scalable MAC engine. iMX in the first preferred embodiment incorporates 4 MAC units; see Figure 19. Alternative preferred embodiments upgrade to 8 MAC units or more. Software

45 can be structured so that the hardware upgrade will not incur substantial software changes. [0222] Much flexibility of iMX is due to parameter-driven address generation and looping control. Overall efficiency comes from efficient pipelining control inside iMX as well as the system-level memory buffering scheme.

[0223] iMX works best for block-based processing. To facilitate this, the datapath needs to connect to data input/output and coefficient memory. iMX contains data input, data output, and coefficient memory ports, and allows arbitration among these ports. This eliminates the need for dedicated memory blocks, and brings more flexibility and better memory utilization on the system level. These memory blocks are accessible as DSP data memory to facilitate data exchange.

50 [0224] There is a separate command memory that feeds a command decode unit in iMX. The command memory should be specified to fit all the accelerated steps in our reference image pipeline algorithm, so that this sequence of commands can be executed with little intervention from DSP.

55 [0225] iMX block diagram appears in Figure 20. A command decode subblock reads and decodes commands, and drives static parameters, one set per command, to the address generator. Address generator then computes looping variables and data/coefficient/output pointers, and coordinates with execution control, which handles cycle-by-cycle

pipelining control. Address generator sends data and coefficient read requests to the arbiter. Arbiter forwards the requests to the data/coefficient memory. Data read back from memory go to the input formatter, which takes care of data alignment and replication. Formatted data and coefficients are then provided to the datapath, which mainly consists of the 4 MAC units. Output from datapath is routed to arbiter for memory write.

[0226] iMX communicates to DSP via shared memory (for data input, coefficient, data output, command) and via memory-mapped registers (start command, completion status). All data buffers and memory blocks are single-ported, and are switched to one party or another via static control, rather than on-line arbitration.

[0227] In a typical application, DSP would place filter coefficients, DCT/IDCT cosine constants, and lookup tables in the coefficient memory, and put iMX commands in the command memory. DSP then turns over access to these memory blocks to iMX. These memory blocks are sized adequately for our reference design to fit all needed coefficients and commands for a major camera operation mode (e.g., image capture). Any update/reload should occur very infrequently. In case either or both memory blocks run out of space, paging can be performed.

[0228] DSP manages the switch network so that, to iMX, there is only one data buffer. During run time, DSP switched the A/B buffers among itself, iMX, VLC, and SDRAM controller to implement data passing.

[0229] Figure 21 illustrates a simple table lookup accelerator with input rounding/clipping capability used to speed up the image pipeline on the DSP. This is carried out with a very simple control structure and datapath.

VLC engine

[0230] VLC accelerator is a coprocessor optimized for quantization and Huffman encode in the context of JPEG compression and MPEG compression. It operates with quantizer matrices and Huffman tables preloaded by DSP, via shared memory blocks. Aggressive pipelining in the design achieves very high throughput rate, above 30 million DCT coefficients for compression.

[0231] VLC's working memory, including quantizer matrices, Huffman tables, and data input/output memory, are all shared memory blocks.

VLC functionality

[0232] Basically, VLC covers Quantization, zigzag scan, and Huffman encode for JPEG encode (baseline DCT, 8-bit sample), with up to 4 quantizer matrices (stored as $\text{invq}[i,j] = 2^{16}/q[i,j]$) and 2 encode Huffman tables all loadable. Can process one MCU that contains up to 10 blocks. Each block consists of $8 \times 8 = 64$ samples.

[0233] Quantization, zigzag scan, and Huffman encode for MPEG-1 video encode. One macroblock, with up to six 8×8 blocks, can be processed. Number of blocks and within them, number of luminance blocks, can be specified. Huffman encode can be bypassed to produce quantized and zigzag-ordered levels.

[0234] The accelerator requires memory blocks for input/output buffer, quantization matrices and Huffman encode tables. The memory configuration should be sufficient to support normal encode operations, one JPEG MCU (minimum coding unit), or MPEG macroblock per call.

[0235] Both input and output must fit the 2K words (lword = 16-bit) shared memory buffer (A or B). MCU or macroblock has maximally ten 8×8 blocks, or 640 input words. Compressed output data is typically smaller than input size.

[0236] JPEG Huffman encode table takes up $(12 \times 176) \times 32\text{-bit}$, or 384 words per table. JPEG standard allows 2 tables, so taking totally 768 memory words. MPEG tables are hard-wired into VLC and do not take up memory. We have allocated 2K words for the Huffman tables.

[0237] The quantizer matrix memory, 512 words by 16-bit, allow for 8 quantizer matrices to coexist, each taking $64 \times 16\text{-bit}$. JPEG allows for 4 matrices, and MPEG encode requires 2 matrices.

Figure 22 shows the major subblocks of VLC. Only the encode path is implemented in one preferred embodiment VLC module; alternative preferred embodiments incorporate the decode path into the module.

ARM subsystem

[0238] ARM microprocessor 130 handles system-level initialization, configuration, user interface, user command execution, connectivity functions, and overall system control. ARM 130 has a larger memory space, better context switching capability, and is thus more suitable for complex, multi-tasking, and general processing than DSP 122. Preferred embodiments integrate an ARM7 cTDMI core; see Figure 1b. ARM7 core is specified up to at least 40 MHz. The ARM subsystem will also have a 32 Kbytes local static RAM 132.

[0239] ARM processor 130 is connected to all the DSC peripherals including CCD Controller, TV encoder, preview engine, IrDA, USB, Compact Flash/Smart Media, UART, etc.

[0240] ARM processor 130 is involved with the management of CCD incoming raw data and intermediate data to the SDRAM and LCD. Connected to all I/O devices, the ARM manages and is responsible for the smart devices such

as USB, IrDA, Compact Flash/Smart Media, and UARTS. The four basic operation modes of PREVIEW, CAPTURE, PLAYBACK, and BURST are initiated by requests from the ARM. The ARM will then monitor the device for completion of the request and in some cases will manage data after the request is completed.

[0241] After RESET and before any of the camera operations can occur, the ARM must perform several housekeeping tasks. The initial task is known as the BOOT operation task. This function not only initializes the I/O and peripherals to a known state, it also must prepare, load and start DSP 122. This sequence begins by reading the DSP boot code from the flash, loading the DSP code memory and then releasing the DSP from its HOLD state. Additional DSP code is loaded into the SDRAM in a format the DSP can then read and overlay into its code space without ARM intervention.

ARM SDRAM Interface

[0242] ARM has two types of access to the SDRAM (1) through SDRAM buffer (burst read/write) and (2) direct access to the SDRAM with a higher latency - 4 cycle READ, 6 cycle WRITE. The direct access to memory can be word, half word or byte access.

[0243] The ARM/SDRAM controller interface also has a 32 byte buffer. The SDRAM burst request first fills this buffer and ARM reads and writes from/to this buffer.

ARM External Memory Interface

[0244] ARM 130 connects to the external memory through the External memory interface module. ARM 130 connects to the Compact Flash/Smart media through this interface. ARM 130 also connects to the off chip flash memory through this interface. DMA block (Figure 1b) enhances the ARM to CF/Smart media transfer.

ARM/DSP BOOT Sequence

[0245] The DSP BOOT sequence begins after a power up or after a COLD START. In this state, DSP 122 is in a HOLD condition waiting on initialization from ARM 130. The ARM checks DSP status registers to assure the DSP is in a HOLD state. The ARM programs the DSP boot code data to the DSP code memory from the FLASH. The code is organized in logical overlays that allow the ARM to select the proper code for the function needed, in this case BOOT code.

[0246] The ARM loads the DSP code using the HPI Bridge (HPIB) interface. This interface can be programmed to access in either 8- or 16-bit width. For BOOT purposes, this will always be a 16bit access.

[0247] After the code is loaded, the ARM signals the DSP to begin by releasing the HOLD. The DSP then begins its reset sequence from an address of DSP 7F80h which is in the DSP RESET vector area. Upon completion of the RESET sequence, the DSP then branches to DSP FF80h, which is the beginning of the BOOT program loaded by the ARM.

[0248] Figure 23a shows the data paths used in the ARM/DSP boot sequence as well as data, request and command exchanges discussed later.

Capture Mode

[0249] ARM 130 programs CCD controller 102 to capture an image. The CCD controller auto transfers the image data to SDRAM and interrupts the ARM using IRQ1 when the transfer is complete. The ARM then notifies the DSP the RAW picture data is available to crunch. When the processing of the raw data is complete, the DSP signals the ARM the task is finished.

Preview Mode

[0250] The CCD will be programmed for a 30 fps high frame rate but reduced resolution vertically. The reconfiguration of the CCD and TG (timing generator) will cause the raw picture data to go to preview engine 104. The DSP will post process the data in SDRAM and prepare parameters for FOCUS, EXPOSURE and WHITE BALANCE. The ARM is signaled by the DSP when new adjustment parameters are ready and those corrections are applied by the ARM. The transferring of the correction parameters use the same communication interrupt architecture as previously mentioned and are expected to be at the current frame rate.

Burst Mode

[0251] The burst mode timing is based on the ARM clocking the picture rate from application parameters. Similar to a cross between Capture and Preview modes, the ARM programs the CCD for a capture that stores a compressed

image into SDRAM through the compression engine. As in Preview mode, the ARM receives adjustment parameters from the DSP to make corrections of FOCUS, EXPOSURE and WHITE BALANCE.

Idle Mode

[0252] ARM may use an idle mode to receive correction parameters from the DSP during periods preceding other camera modes. If not in a power down situation, this time of 10-15 frames will allow the DSP-to-ARM correction loop to make auto corrections on FOCUS, EXPOSURE and WHITE BALANCE. This idle mode will simulate Preview mode for the purposes of obtaining a stable correction.

ARM/DSP communication

[0253] The communication between ARM 130 and DSP 122 is via the HPIB (Host Port Interface Bridge). The HPIB physically connects the DSP (a C5409 type DSP) ports and BUSC (BUS Controller) 134. The ARM accesses the DSP memory by programming the HPIB, opening a 32k-word window into the DSP memory map. The map contains the data structures shared by the ARM and DSP for command request's, acknowledgements and datagrams.

[0254] The HPIB contains five sub-blocks. They are the interface, timing generator, DSP control registers, and interrupt hold sections.

[0255] The interface section receives and stores data from BUSC 134 and transfers it to and from the C5409. This interface can be an 8- or 16-bit data path to the C5409 and is 16-bit to the BUSC. An added feature is the ability to exchange the upper and lower byte if programmed to do so.

[0256] The timing generator makes signals HBIL and HDS and detects signal HRDY. HBIL is the HPI byte identification signal- to the C5409. The HDS is the data strobe signal to the C5409 and the HRDY is the ready signal read from the C5409.

[0257] The interrupt hold section will detect the HINT level and make the INTC pulse synchronized with the ARM clock. The module will also set the HOLD port of the C5409 and detect HOLDA.

[0258] In 8-bit mode, address data from the ARM will not reach the C5409. The address is used only if the C5409 internal memory is selected. Therefore, the ARM must set the address in the HPIA register before sending or receiving data to the 32 Kword DARAM. The 8-bit mode may also be used for ARM<->DSP handshaking. The ARM will use the HINT bit in the HPIC register to interrupt the C5409.

[0259] In 16-bit mode, the HPIA/HPIC/HPID are not used. The ARM can access the C5409 internal memory as if it exists in the HPIB module. This mode will deliver faster performance, but does not support the HANDSHAKE signals because of these are routed in the HPIC register.

[0260] Figure 23b shows the signals and paths for the ARM to reach the C5409 DARAM.

[0261] Figure 23c indicates the shared memory map between the ARM (HOST) and the C5409 processor. When the ARM selects the memory area, "DSP Memory", BUSC takes cs_hpb signal active. The ARM can now access the DSP internal memory (32 kword DARAM + HPIA + HPIC + HPID).

[0262] When the ARM selects the "DSP Controller" area, BUSC takes cs_dspsc signal active. The ARM is now accessing registers related to the C5409.

Multi-processing debugging environment

[0263] The preferred embodiment integrates ARM 130 and DSP 122 and thus multi-processing and thus requires debugging and development support. The preferred embodiment accomplishes this with a single JTAG connector 170 with additional emulation logic as illustrated in Figure 24.

Input/Output modules

[0264] The input/output module provides the different interfaces with the DSC peripherals as follows.

[0265] TV encoder 106 produces NTSC/PAL and RGB outputs for the LCD display and TV.

[0266] CCD/CMOS controller 102 generates timing signals VD/HD, can synchronize on externally generated HD/VD signals (#0 of MODESET register, #0 of SYNCEN register), supports progressive scan and interlaced CCDs, generates black clamping control signals, programmable culling pattern 9CULH, CULV registers), 1 line/2 line alternating fields, MCLK (generated by CCD module), WEN (WRQ on TG, active-high) indicates CCD controller writing data to SDRAM, TG serial port interface (clk, data, TG chip select) is controlled by GIO pins, Iris, mechanical shutter, focus and zoom are controlled by GIO pins.

[0267] USB 142 from programmer's perspective consists of three main parts: FIFO controllers, UDC controller, and UDC core. USB configuration: INTERFACED0 ALT0 ENDPOINT0: CONTROL; INTERFACE0 ALT0 ENDPOINT1:

BULKIN; INTERFACE0 ALT0 ENDPOINT1: BULKOUT; INTERFACE1 ALT0 ENDPOINT2: ISOIN; INTERFACE2 ALT0 ENDPOINT3: INTERRUPT IN. Buffer configuration: SUB module has six FIFOs inside; each FIFO is of the same construction, except for direction and buffer size: USB module has only one unified memory for all endpoints; buffer sizes are programmable as long as all buffers fit inside the memory.

[0268] UART part of I/O block 140, supports start/stop communication protocol, detects parity errors (supporting data length of 7 or 8 bits with even, odd, or no parity and 1 or 2 stop bits), has 32 bytes of FIFO for both transmitter and receiver, generates interrupts for a FIFO overflow or a time-out is detected on data receiving. ARM 130 control UART modules. There are seven 16-bit width registers which are accessible from ARM 130: data transmitter/receiver register (FIFO), bit rate register, mode register, FIFO control register for receiver, FIFO control register for transmitter, line control register, and status register. Figure 25 is a block diagram.

[0269] Compact Flash/Smart Media interface 180 is used to save/store image or user's data to a compact flash card or smart media; see Figure 26. The interface supports two kinds of operation modes for register setting and data transfer: memory mapped mode and I/O mode. An ARM 130 interrupt is generated for card detection while a compact flash card is being plugged or unplugged. The pins for both the smart media and the compact flash control interfaces are overlapped and can be switched by ARM 130 depending on product needs; see Figure 26.

[0270] In particular, the compact flash controller has registers mapped to the ARM memory space. The compact flash controller is responsible for generating the related control signals to the interface pins, and writes at 420 KB/s and reads at 2.0 MB/s. SDRAM can be utilized for storing at least one picture and an attempt to write to the compact flash with a big sector count, as done in a DOS machine, will invoke the fast write performance.

[0271] In contrast, the smart media controller has five register settings: command register, address1 register, address2 register, address3 register, and data port register. These five registers are mapped to the ARM memory space, and smart media controller will generate the related signals for different register access automatically.

[0272] Audio input/output may be through the serial port of I/O block 140 with DSP buffering.

[0273] Infrared data access (IrDA) is supported by a fast FIR core and part of I/O block 140.

[0274] Block 140 also contains general purpose input/output which can support items such as CCD/CMOS imager module control for tuning AGC gain and electronic shutter, RTC control, battery power detection which can generate inner interrupt to the ARM for appropriate system response, camera lens motor control for focus and zoom, a user keypad input, LED indicators, flash light control, and power management control.

iMX programming

[0275] DSP 122 instructs iMX 124 to perform tasks by sending iMX commands. These commands can be complex to understand and contain many parameters that are fixed in the inner loops. The ideal model is to provide separate command building and command transfer routines to the DSP programmer, so that the commands can be pre-constructed outside the loop, and transferred to iMX as generic data memory moves inside the loop. Commonly used iMX commands are prepackaged in C code to ease the programming.

ARM/DSP task allocation

[0276] ARM 130 runs an operating system such as Windows CE, controls low frequency, synchronous input/output (such as to a compact flash card (CFC), and controls user interactions which also are slow and all the peripheral modules control preview engine, burst mode compression, TV encoder, CCD controller, USB, CF, IrDA, etc.

[0277] DSP 122 runs an operating system such as SPOX, controls all real-time functions (auto focus, auto exposure, auto white balance), real-time input/output (audio IO, modem IO), real-time applications (e.g., audio player), computational expensive signal processing tasks (image pipeline, JPEG 2000, image stitching).

Pin description of integrated circuit chip

[0278] The preferred embodiment pins are as follows

CCD SENSOR	
Pin Count : 16	
1.C_PCLK	(I) Pixel clock
2.C_VSYNC	(I/O) Vertical sync
3.C_HSYNC	(I/O) Horizontal sync
4.C_FIELD	(I/O) Field indicator

(continued)

CCD SENSOR	
Pin Count : 16	
5. C_WEN	(I) CCDC write enable
6:17.C_DATA	(I) Image data 12Bit

SDRAM Interface	
Pin Count : 58	
1.SDR_CLK	(O) Master clock
2.SDR_CKE	(O) Clock enable
3.SDR_WE	(O) Write enable
4.SDR_CAS	(O) Column address strobe
5.SDR_RAS	(O) Row address strobe
6.SDR_CS0	(O) Support 2pc of RAM
7.SDR_CS1	(O) Support 4pc of RAM
8:39.DQ[31:0]	(I/O) Data bus
40:54.SDR_A[14:0]	(O) Address bus
55.SDR_DQMHH	(O) DQM for DQ[31:24]
56.SDR_DQMH	(O) DQM for DQ[23:16]
57.SDR_DQMLH	(O) DQM for DQ[15:8]
58.SDR_DQMLL	(O) DQM for DQ[7:0]

ARM BUS	
Pin Count : 39	
1:23.ARM_A[22:0]	(O) Address bus
24:39.ARM_D[15:0]	(O) Data bus

Audio Interface	
Pin Count : 6	
1.DSP_BDX	(O) Serial port transmit
2.DSP_BCLKX	(I/Q) Transmit clock
3.DSP_BFSX	(I/O) Frame synchronization pulse
4.DSP_BDR	(I) Serial data receive
5.DSP_BCLKR	(I) Receive clock
6.DSP_BFSR receive	(I) Frame synchronization pulse

External Flash Interface	
Pin Count : 5	
1.FLSH_WE	(O) Write enable
2.FLSH_CE	(O) Chip select
3.FLSH_OE	(O) Output enable
4.FLSH_SIZE	(I) 8Bit/16Bit select
5.FLSH_BSY	(I) Busy input

5

10

USB(T.B.D)	
Pin Count : 10	
1.M48XO	(O) 48MHz clock output
2.M48XI	(I) 48MHz clock input
3.USB_DP	(I/O) Differential data+
4.USB_DM	(I/O) Differential data-
5.ATTACH	(I) Attach detect

15

20

UART	
Pin Count : 5	
1.RXD	(I) UART RX
2.TXD	(O) UART TX
3.ERXD	(I) UART Rx for external CPU
4.ETXD	(O) UART Tx for external CPU
5.SIFDO	(O) Serial I/F data output

25

IrDA	
Pin Count : 2	
1.IRXD	(I) IrDA RX
2.ITXD	(O) IrDA TX

30

35

40

Compact Flash	
Pin Count : 9	
1.CFE1	(O) Card enable#1
2.CFE2	(O) Card enable#2
3.IOIS16	(O) I/O select
4.STSCHG	(I/O) Status changed
5.CFWAIT	(I) Wait signal input
6.CFRST	(O) Reset
7.CFD1	(I) Card Detect pin#1
8.CFD2	(I) Card Detect pin#2
9.CFRDY	(I) Ready

45

50

55

TV/RGB DAC Analog output	
Pin Count : 27	
1.IREF(R)	(I) R-ch Current reference control
2.DAOUT(R)	(O) Analog output R-ch
3.GNDA	Analog GND
4.VCCA	Analog VCC
5.BIAS	(I) Phase compensation cap.R-ch
6.VREF	(I) RGB common reference voltage
7.IREF(G)	(I) G-ch Current reference control
8.DAOUT(G)	(O) Analog output G-ch
9.GNDA	Analog GND

(continued)

TV/RGB DAC Analog output	
Pin Count : 27	
10. VCCA	Analog VCC
11. BIAS	(I) Phase compensation cap.G-ch
12. IREF(B)	(I) B-ch Current reference control
13. DAOUT(B)	(O) Analog output B-ch
14. GNDA	Analog GND
15. VCCA	Analog VCC
16. BIAS	(I) Phase compensation cap.B-ch
17. IREF(C)	(I) Composite Current reference control
18. DAOUT(C)	(O) Analog output Composite
19. GNDA	Analog GND
20. VCCA	Analog VCC
21. VREF	(I) Composite reference voltage
22. BIAS	(I) Phase compensation cap.composite
23. DVCC	Digital VCC for DAC
24. DGND	Digital GND for DAC
25. HSYNC	(O) H-sync output for RGB output
26. VCSYNC	(O) V-sync / Composite-sync(select by register)

GIO	
Pin Count : 32 [31:0]	
1:32.GIO	(I/O) General Purpose I/O

Miscellaneous	
Pin Count : 15	
1.RESET	(I) Power on reset
2.M27XI	(I) 27MHz input
3.M27XO	(O) 27MHz output
4.TCK	(I) JTAG clock
5.TDI	(I) JTAG data input
6.TDO	(O) JTAG data output
7.TMS	(I) JTAG test mode select
8.TRST	(I) JTAG test reset
9.EMU0	(I/O) Emulator interrupt 0 pin
10.EMU1	(I/O) Emulator interrupt 1 pin
11.TESTO	(I) Test input 0
12.TEST1	(I) Test input 1
13.SCAN	(I) Test input
14.TESTSL0	(I) Test mode select 0
15.TESTSL1	(I) Test mode select 1

TOTAL PIN COUNT		
CCD SENSOR		17
SDRAM I/F	58	

(continued)

TOTAL PIN COUNT		
ARM BUS	39	
Audio I/F	6	
Flash memory I/F		5
USB	5	
UART	5	
IrDA	2	
Compact Flash I/F		9
4DAC	26	
GIO	32	
Miscellaneous		15
Sub Total		219 pins
Power :	37 pins (14%)	
TOTAL	: 256 pins	
Audio player		

[0279] Portable digital audio players are expected to be one of the most popular consumer products. Currently the MP-3 player based on MPEG-1 Layer 3 audio compression standard is growing rapidly in portable audio market while MPEG-2 AAC and Doby AC-3 are alternative digital audio coding formats to be considered as emerging standards. Thus the preferred embodiments' programmability permits inclusion of digital audio player functions. The audio can be input via flash memory, PC, etc. and the decoded can be output on the serial port. The decoding program can be loaded from flash memory, ROM, etc.

[0280] Insofar as embodiments of the invention described above are implementable, at least in part, using a software-controlled programmable processing device such as a Digital Signal Processor, microprocessor, other processing devices, data processing apparatus or computer system, it will be appreciated that a computer program for configuring a programmable device, apparatus or system to implement the foregoing described methods is envisaged as an aspect of the present invention. The computer program may be embodied as source code and undergo compilation for implementation on a processing device, apparatus or system, or may be embodied as object code, for example. The skilled person would readily understand that the term computer in its most general sense encompasses programmable devices such as referred to above, and data processing apparatus and computer systems.

[0281] Suitably, the computer program is stored on a carrier medium in machine or device readable form, for example in solidstate memory or magnetic memory such as disc or tape and the processing device utilises the program or a part thereof to configure it for operation. The computer program may be supplied from a remote source embodied in a communications medium such as an electronic signal, radio frequency carrier wave or optical carrier wave. Such carrier media are also envisaged as aspects of the present invention.

[0282] In view of the foregoing description it will be evident to a person skilled in the art that various modifications may be made within the scope of the invention.

[0283] The scope of the present disclosure includes any novel feature or combination of features disclosed therein either explicitly or implicitly or any generalisation thereof irrespective of whether or not it relates to the claimed invention or mitigates any or all of the problems addressed by the present invention. The applicant hereby gives notice that new claims may be formulated to such features during the prosecution of this application or of any such further application derived therefrom. In particular, with reference to the appended claims, features from dependent claims may be combined with those of the independent claims and features from respective independent claims may be combined in any appropriate manner and not merely in the specific combinations enumerated in the claims.

Claims

1. An integrated circuit for a digital still camera, comprising:

- (a) a first programmable processor programmed to run control functions, said first processor coupled to a user interface, a controller for memory, and a controller for image acquisition;
- (b) a second programmable processor programmed to run image processing and compression functions, said second processor coupled to said first processor; and

(a) a third processor coupled to said second processor, said third processor including at least four parallel multiply and accumulate units.

2. A method of resizing a digital image, comprising the steps of:

- (a) providing an M1 by M2 pixel image where M1 and M2 are positive integers;
- (b) filtering a row of M1 pixels of the image with a set of filter coefficients which depends upon pixel location to yield M1 filtered outputs;
- (c) discarding M1-N1 of the filtered outputs of step (b);
- (d) repeating steps (b) and (c) for all other rows of the M1 by M2 image to yield a N1 by M2 image;
- (e) filtering a column of M2 pixels of the N1 by M2 image with a set of filter coefficients which depends upon pixel location to yield M2 filtered outputs;
- (f) discarding M2-N2 of the filtered outputs of step (e); and
- (g) repeating steps (e) and (f) for all other rows of the N1 by M2 image to yield a N1 by N2 image.

3. A method of color-filtered array interpolation, comprising the steps of:

- interpolate a first color subarray to yield a first color array;
- high-pass filter the first color array from the prior step to yield a high pass first color array;
- interpolate a second color subarray to yield a second color array;
- weight and add the high-pass first color array to the second color array to yield a modified second color array;
- interpolate a third color subarray to yield a third color array;
- weight and add the high-pass first color array to the third color array to yield a modified third color array;
- defining a three-color image with the first color array, the modified second color array, and the modified third color array.

4. A method of tone-scaling an image, comprising the steps of:

- (a) providing an M by N pixel image where M and N are positive integers and each pixel has an intensity in the range from 0 to K;
- (b) determining a scaled cumulative distribution function $T(X)$ for X in the range 0 to K and with values in the range 0 to K. $T(X)$ equals K multiplied by the number of pixels with intensity less than or equal to X and then divided by MN;
- (c) selecting a parameter A in the range 0 to 1;
- (d) adjusting an intensity Z of a pixel of step (a) to $AT(Z)+(1-A)Z$.

5. A method of interpolation for a Bayer color-filtered array, comprising the steps of:

- interpolate the green subarray to form a tentative green array;
- clamp the interpolated pixel values of the tentative green array to lie in the range of the middle two values of the four neighboring values of the green subarray to form a green array; and
- interpolate the red and blue subarrays.

6. A method of interpolation for a Bayer color-filtered array image, comprising the steps of:

- interpolate the green subarray to form a green array;
- diagonally interpolate the red subarray to form a second red subarray;
- horizontally and vertically interpolate the second red subarray to form a red array;
- diagonally interpolate the blue subarray to form a second blue subarray;
- horizontally and vertically interpolate the second blue subarray to form a blue array;
- wherein the diagonal interpolation includes weightings by the values of the corresponding pixels in the green array;
- wherein the horizontal and vertical interpolation includes weightings by the values of the corresponding pixels in the green array; and
- the green array, red array, and blue array define a color image.

7. A method of white balancing for a color image sensor, comprising the steps of:

- (a) providing output green, red, and blue pixel intensities for a color image sensor illuminated with white light;
- (b) determining the ratio of the green pixel intensities to the blue pixel intensities for the outputs of step (a);
- (c) providing an adjustment for blue pixel intensities as multiplication by the ratio of step (b) plus addition of an offset which depends upon the pixel intensities in step (a);
- (d) determining the ratio of the green pixel intensities to the red pixel intensities for the outputs of step (a); and
- (e) providing an adjustment for red pixel intensities as multiplication by the ratio of step (b) plus addition of an offset which depends upon the pixel intensities in step (a).

8. A method of video decoding and display, comprising the steps of:

- (a) providing a sequence of encoded video frames, a sequence of video frame boundary times, and storage locations for decoded ones of the frames;
- (b) when the m'th frame of the sequence has been decoded, next decode the m'th frame of the sequence with m' determined as TP_n divided by DeltaT where DeltaT is the time interval between successive ones of the video frame boundary times of step (a) and TP_n is a presentation time determined as the greater of (i) the presentation time TP_{n-1} of the m'th frame plus DeltaT and (ii) the starting time T_s for decoding the frame to be presented next after the m'th frame plus an estimated time to decode a frame, the sum rounded up to a multiple of DeltaT, wherein the starting time T_s is the greater of (i) current time and (ii) the time a storage location to contained the decoded m'th frame will be available due to a prior decoded frame being displayed;
- (c) displaying a decoded frame;
- (d) repeating steps (b) and (c) with incremented indices.

9. A method of interpolation for a complementary-color-filtered array image, comprising the steps of:

- (a) provide a complementary-color-filtered array of pixel values with yellow pixel values Y_e on a first subarray, cyan pixel values C_y on a second subarray, magenta pixel values M_g on a third subarray, and green pixel values G on a fourth subarray;
- (b) interpolating the yellow pixel values to all pixels in the array;
- (c) interpolating the cyan pixel values to all pixels in the
- (d) interpolating the magenta pixel values to all pixels in the array;
- (e) interpolating the green pixel values to all pixels in the array;
- (f) adjusting the pixel yellow value by subtracting quantity $(Y_e + C_y - 2 \cdot G - M_g)/4$ from the pixel yellow value from step (b) where Y_e is the pixel yellow value from step (b), C_y is the pixel cyan value from step (c), M_g is the pixel magenta value from step (d), and G is the pixel green value from step (e);
- (g) adjusting the pixel cyan value by subtracting quantity $(Y_e + C_y - 2 \cdot G - M_g)/4$ from the pixel cyan value from step (c) where Y_e is the pixel yellow value from step (b), C_y is the pixel cyan value from step (c), M_g is the pixel magenta value from step (d), and G is the pixel green value from step (e);
- (h) adjusting the pixel magenta value by adding quantity $(Y_e + C_y - 2 \cdot G - M_g)/4$ to the pixel magenta value from step (d) where Y_e is the pixel yellow value from step (b), C_y is the pixel cyan value from step (c), M_g is the pixel magenta value from step (d), and G is the pixel green value from step (e); and
- (i) adjusting the pixel green value by adding quantity $(Y_e + C_y - 2 \cdot G - M_g)/8$ to the pixel green value from step (e) where Y_e is the pixel yellow value from step (b), C_y is the pixel cyan value from step (c), M_g is the pixel magenta value from step (d), and G is the pixel green value from step (e).

10. A method of interpolation for a color-filtered array image, comprising the steps of:

- (a) provide a color-filtered array of pixel values with red pixel value $R(j,k)$ for the pixel located at the jth row and kth column when j and k are even integers, blue pixel value $B(j,k)$ for pixel located at jth row and kth column when j and k are odd integers, and green pixel value $G(j,k)$ for the pixel located at the jth row and kth column when one of j and k is an even integer and the other an odd integer;
- (b) horizontally interpolate rows of pixel values when j is even to yield rows of red values $R(j,k)$ and green values $G(j,k)$ for all k, and when j is odd to yield green values $G(j,k)$ and blue values $B(j,k)$ for all k;
- (c) when j is an even integer vertically interpolate the horizontal interpolated red rows from step (b) to yield the red values $R(j-1,k)$ and when j is an odd integer vertically interpolate the horizontal interpolated blue to yield the blue values $B(j-1,k)$;
- (d) vertically filter the horizontal interpolated green pixel rows from step (b) to yield filtered green output values $G^*(j,k)$;
- (e) define the difference values $D(j,k)$ between the green values $G(j,k)$ and the filtered green values $G^*(j,k)$;

(f) when j is even subtract $D(j,k)$ from the red values $R(j,k)$ to yield red output values $R''(j,k)$ and when m is odd add $D(j,k)$ to the red values $R(j,k)$ to yield red output values $R''(j,k)$; and
 (g) when m is even add $D(j,k)$ to the blue values $B(j,k)$ to yield blue output values $B''(j,k)$ and when m is odd subtract $D(j,k)$ from the blue values $B(j,k)$ to yield blue output values $B''(j,k)$.

5

11. A method of playback buffering, comprising:

- (a) providing first and second buffers, each of the buffers with a full/not-full indicator, and a position indicator indicating a position within the first and second buffers;
- 10 (b) filling the first buffer with encoded bits when the full/not-full indicator for the first buffer indicates not-full and the position indicator indicates a position in the second buffer;
- (c) filling the second buffer with encoded bits when the full/not-full indicator for the second buffer indicates not-full and the position indicator indicates a position in the first buffer;
- 15 (d) decoding a group of bits in the first and second buffers starting at the position indicated by the position indicator and treating the first and second buffers as forming a single circular buffer;
- (e) updating the full/not-full indicators according to the bits removed by step (d); and
- (f) repeating steps (b)-(e).

20

12. A computer program comprising computer- or machinereadable computer program elements for configuring a computer to implement the method of any one of claims 2 to 11.

13. A computer program comprising computer- or machinereadable computer program elements translatable for configuring a computer to implement the method of any one of claims 2 to 11.

25

14. A carrier medium carrying a computer program according to claim 12 or 13.

30

35

40

45

50

55




FIG. 1a

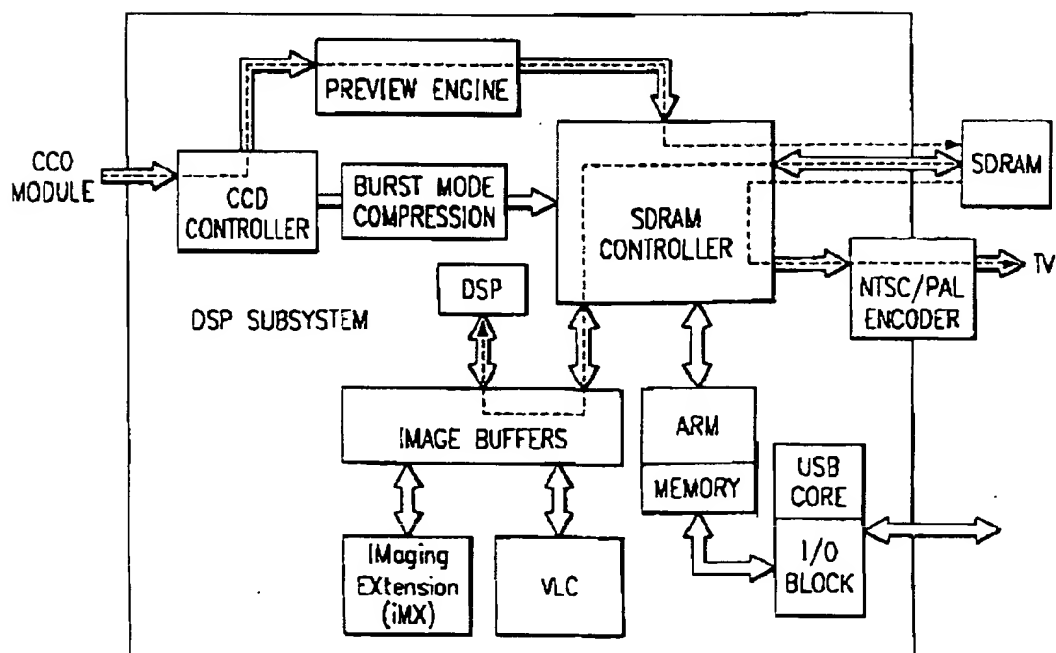



FIG. 2



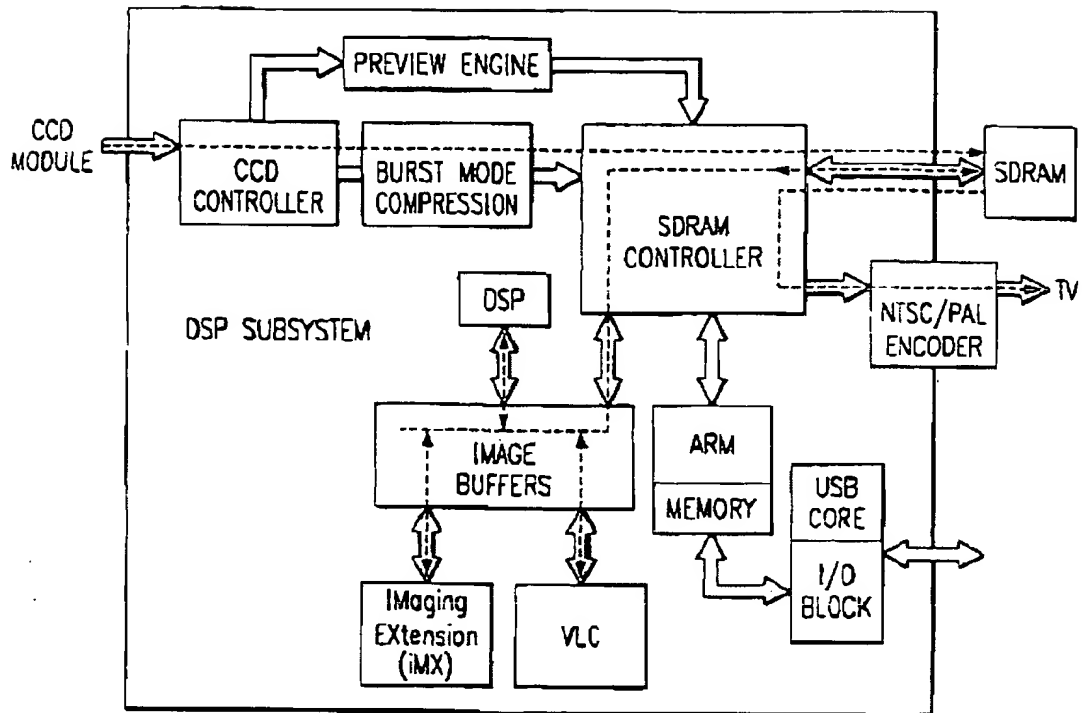


FIG. 3a

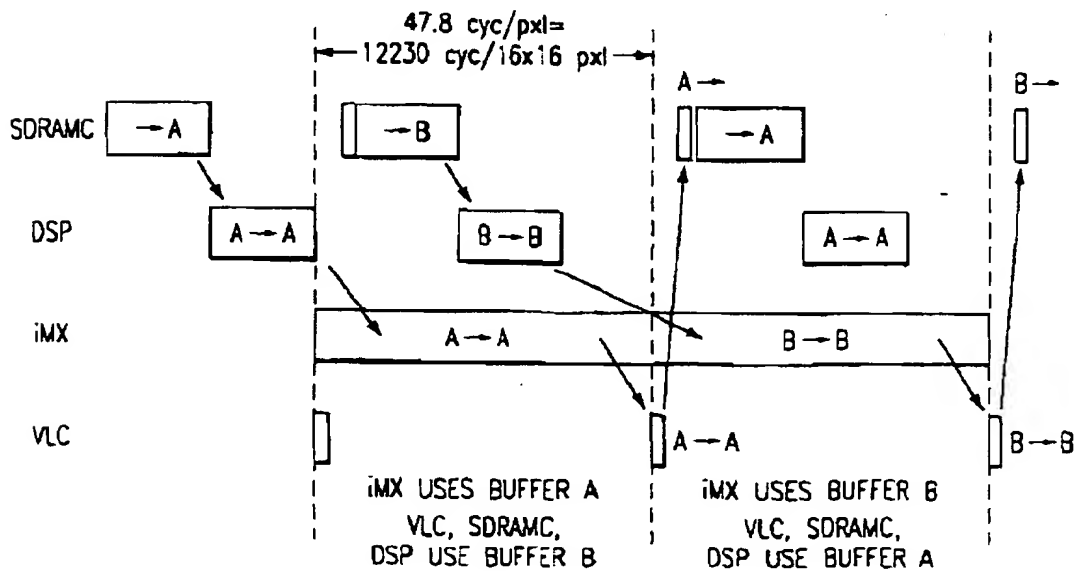


FIG. 3b




FIG. 4

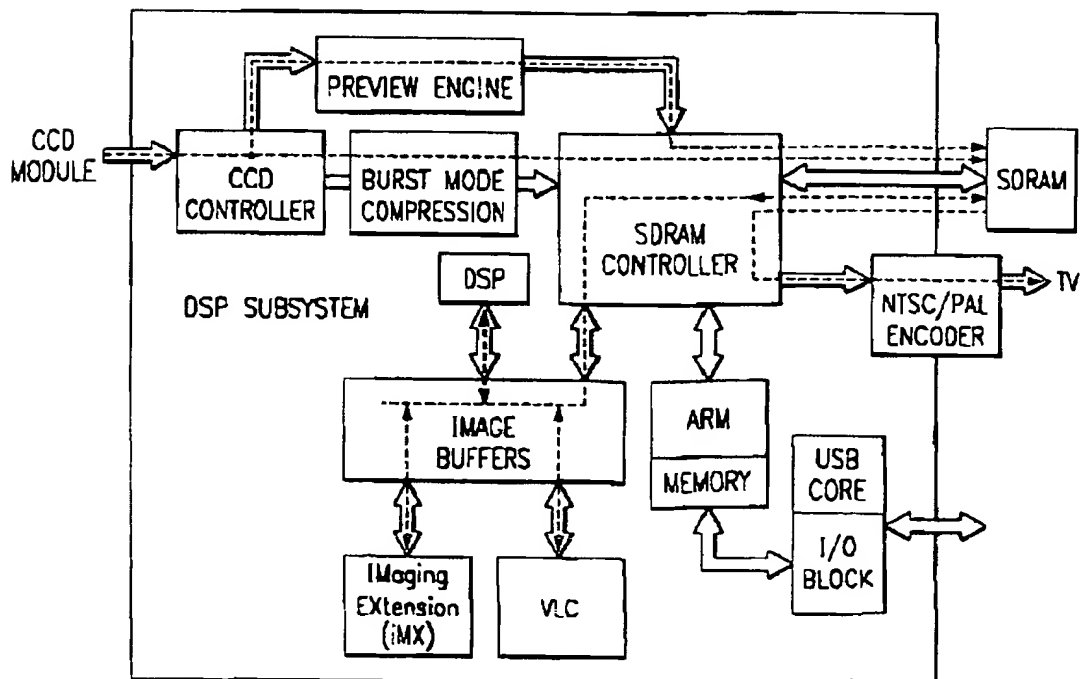


FIG. 5

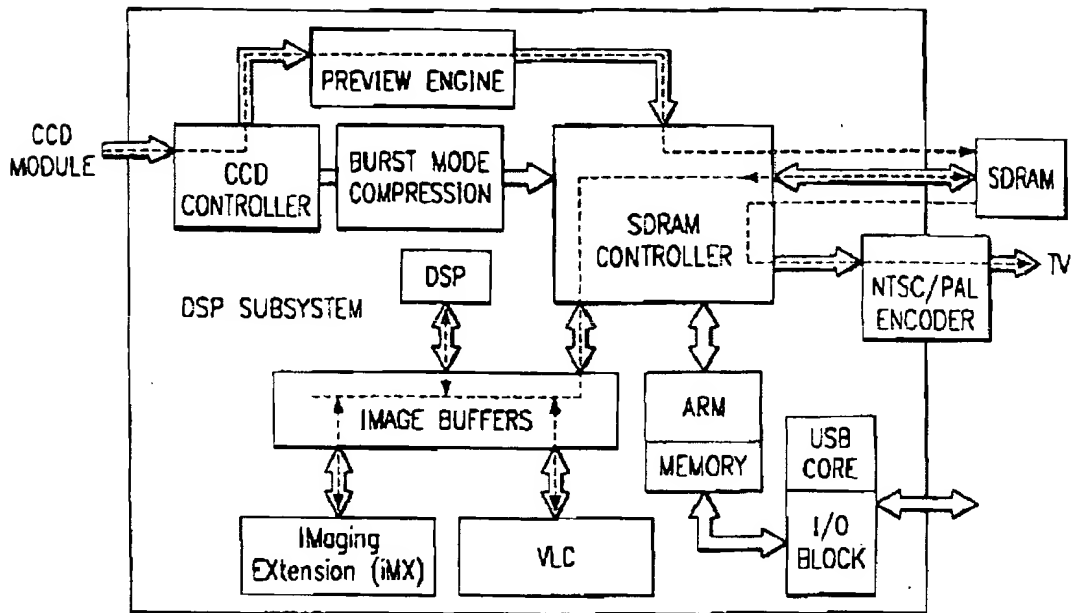


FIG. 6

R	G	R	G
G	B	G	B
R	G	R	G
G	B	G	B

FIG. 7a

Ye	Cy	Ye	Cy
G	Mg	G	Mg
Ye	Cy	Ye	Cy
G	Mg	G	Mg

FIG. 7b

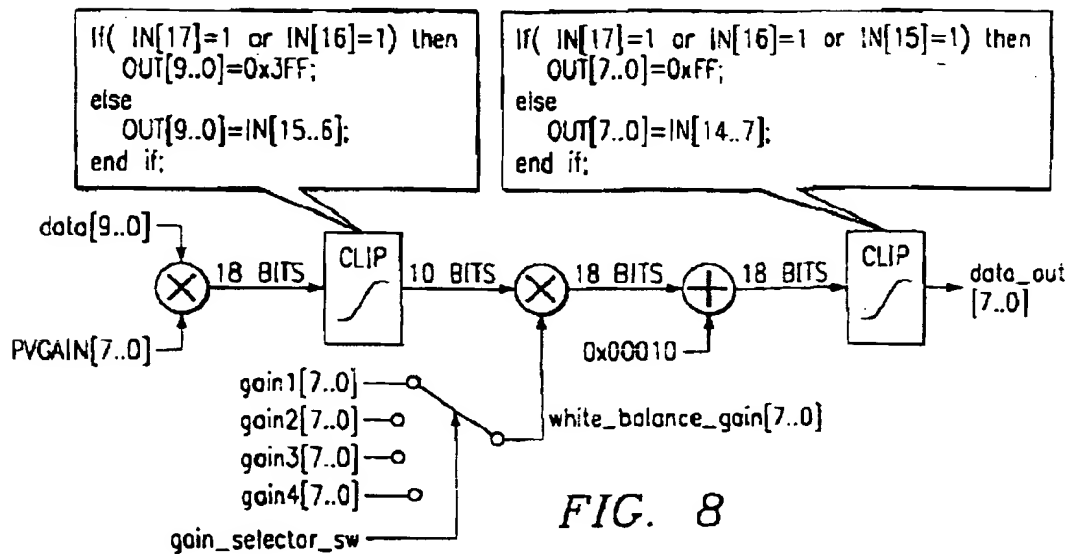


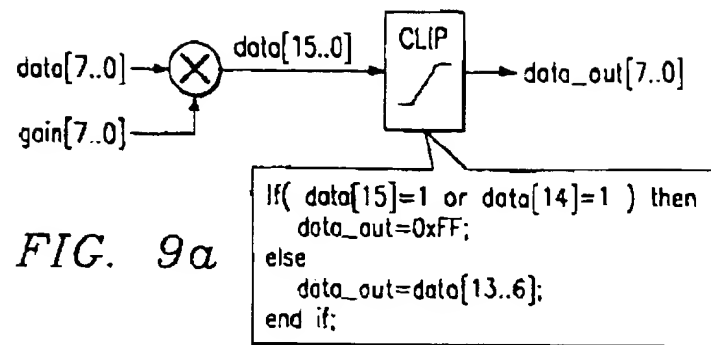


FIG. 8






FIG. 10a


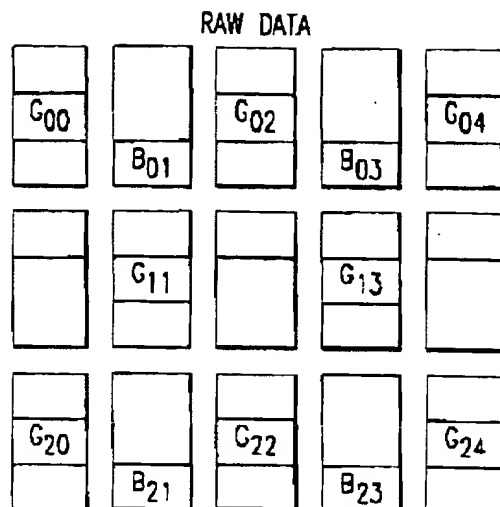



FIG. 10b



HORIZONTAL
INTERPOLATION



NORMAL MODE

$$g_{12} = \frac{-R_{10} + 2G_{11} + 2R_{12} + 2G_{13} - R_{14}}{4}$$

SIMPLE MODE

$$g_{12} = \frac{G_{11} + G_{13}}{2}$$

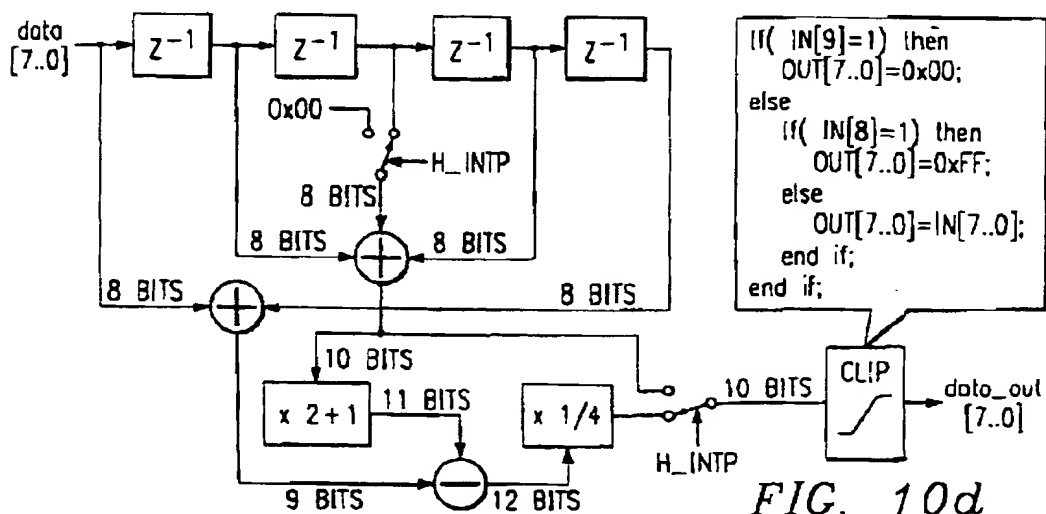
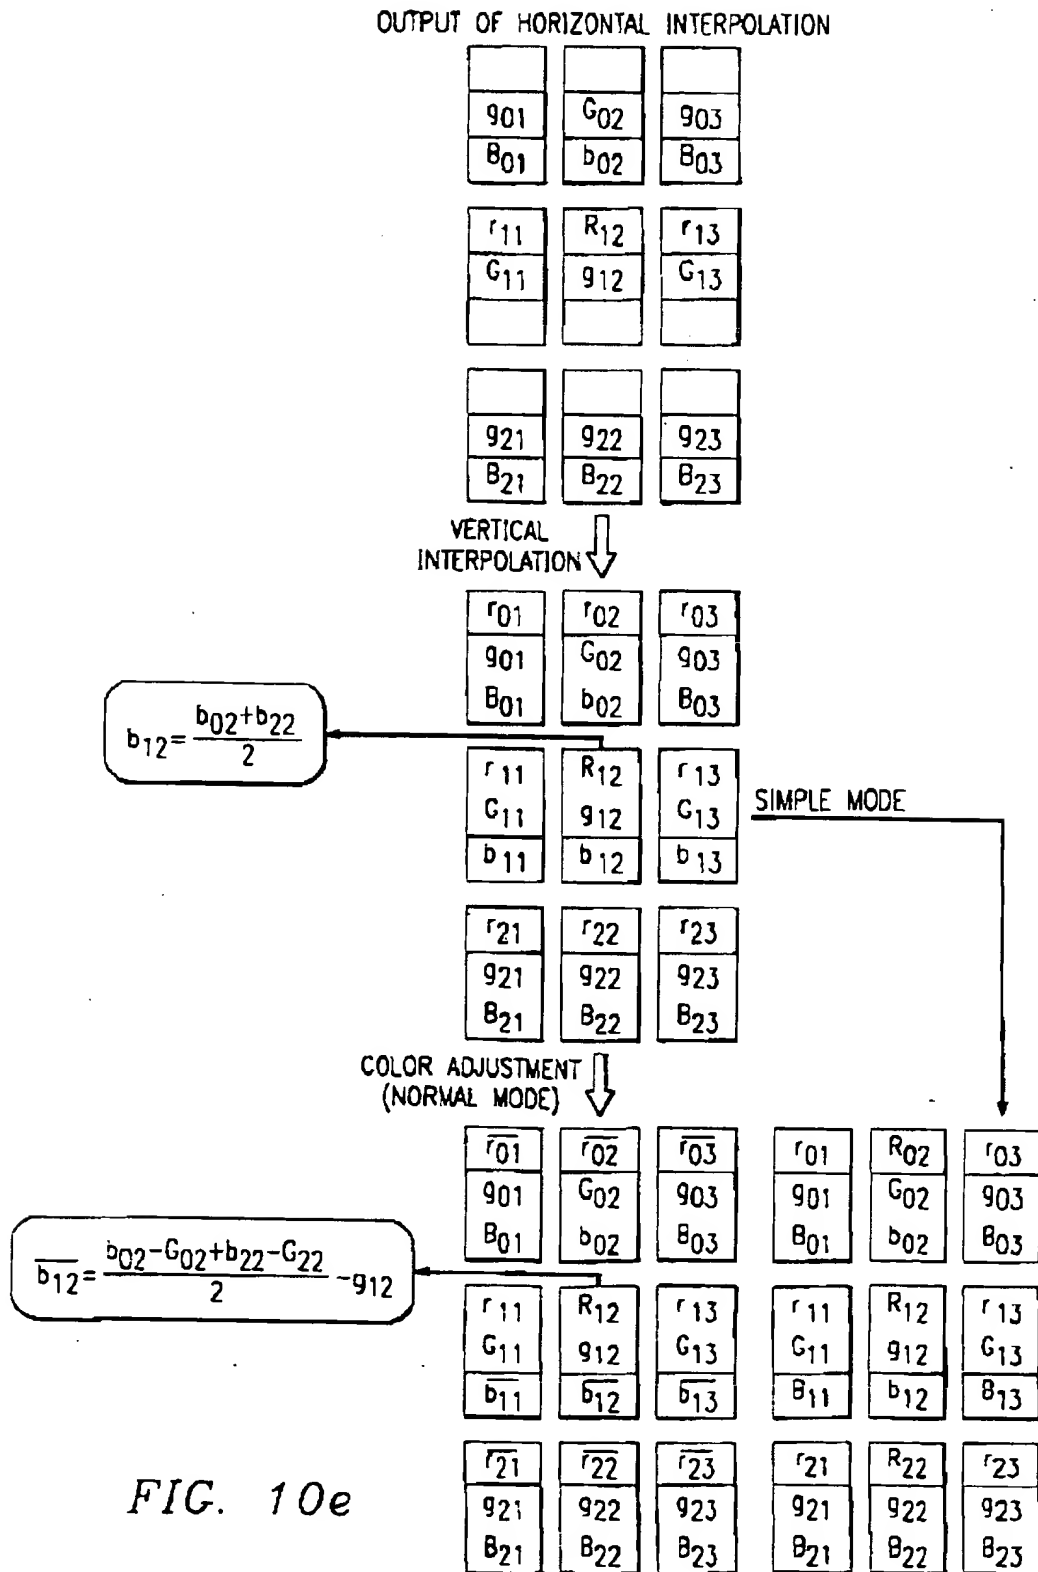


FIG. 10d



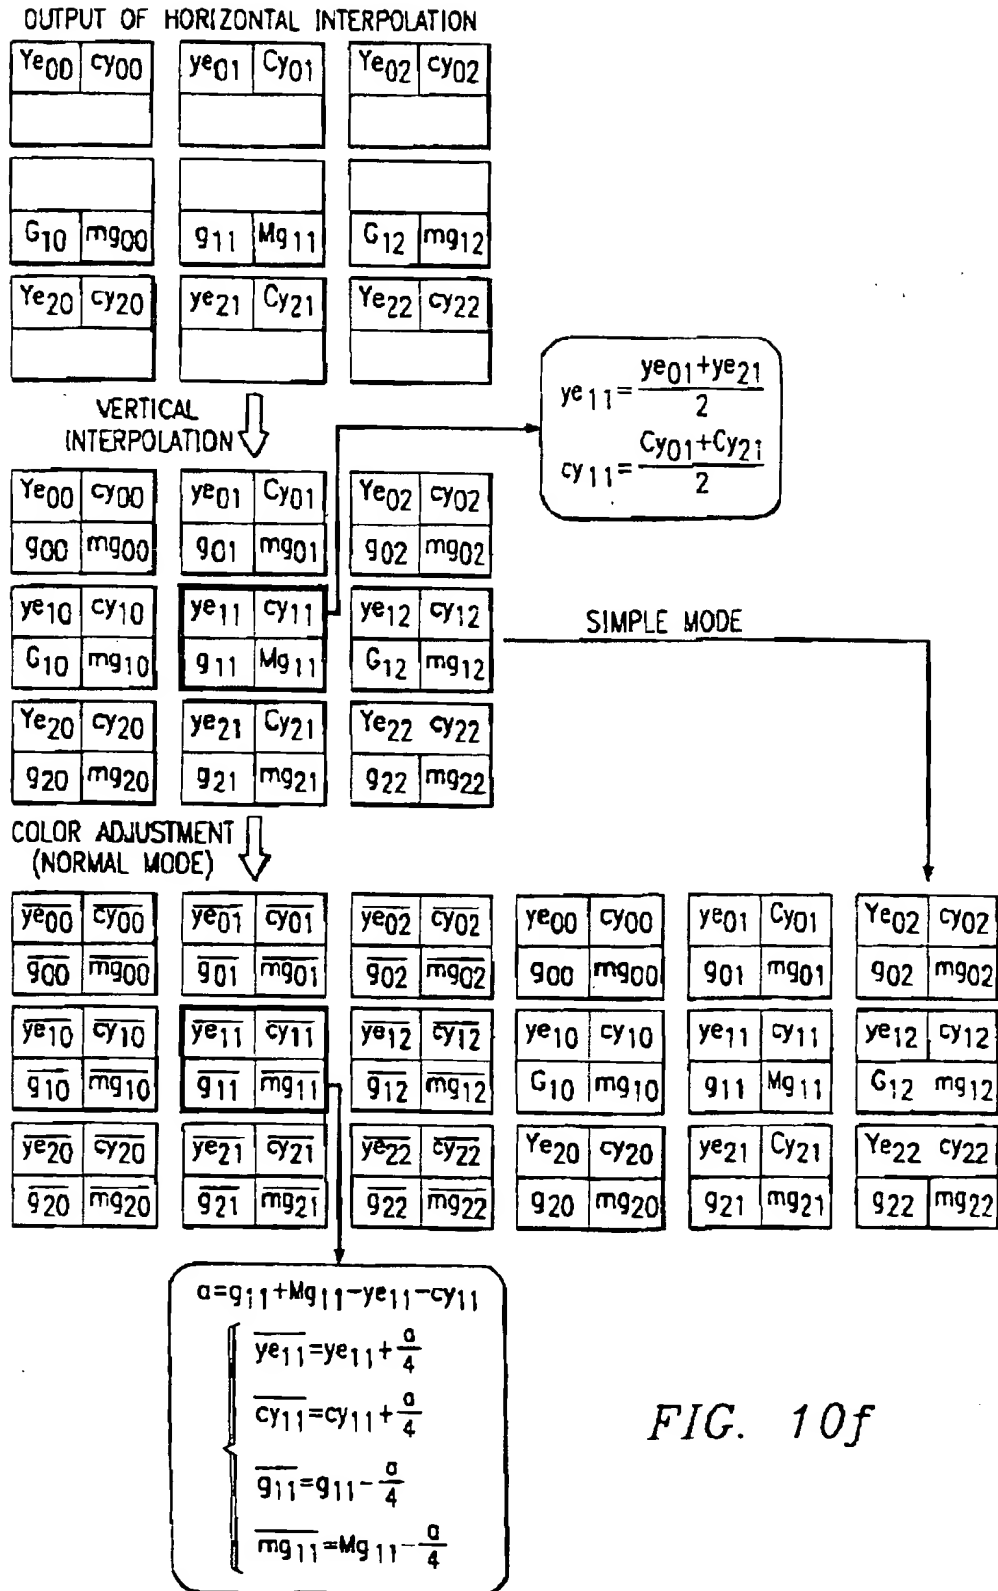
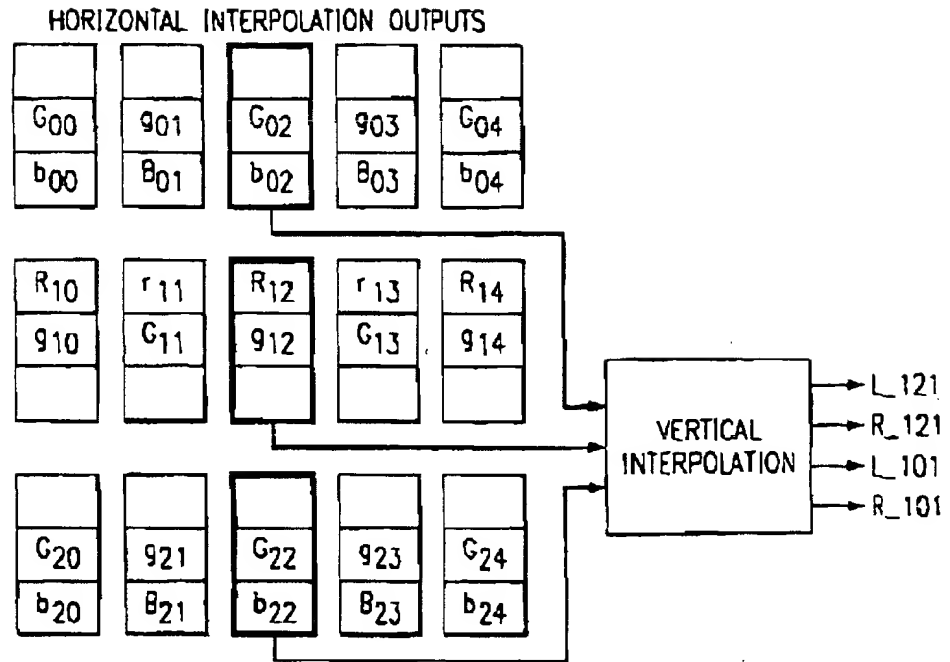


FIG. 10f



NOISE FILTER = OFF

$$\left\{ \begin{array}{l} L_{121} = R_{12} \\ R_{121} = g_{12} \\ L_{101} = \frac{G_{02} + G_{22}}{2} \\ R_{101} = \frac{b_{02} + b_{22}}{2} \end{array} \right.$$

NOISE FILTER = ON

$$\left\{ \begin{array}{l} L_{121} = R_{12} - g_{12} + \frac{G_{02} + 2g_{12} + G_{22}}{4} \\ R_{121} = \frac{G_{02} + 2g_{12} + G_{22}}{4} \\ L_{101} = \frac{G_{02} + G_{22}}{2} \\ R_{101} = \frac{b_{02} + b_{22}}{2} \end{array} \right.$$

FIG. 10g




FIG. 10h

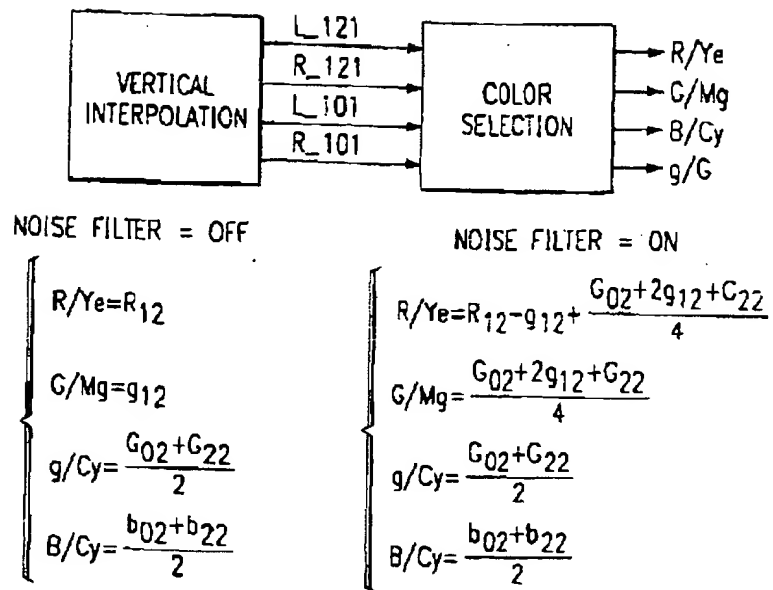


FIG. 10i




FIG. 10j

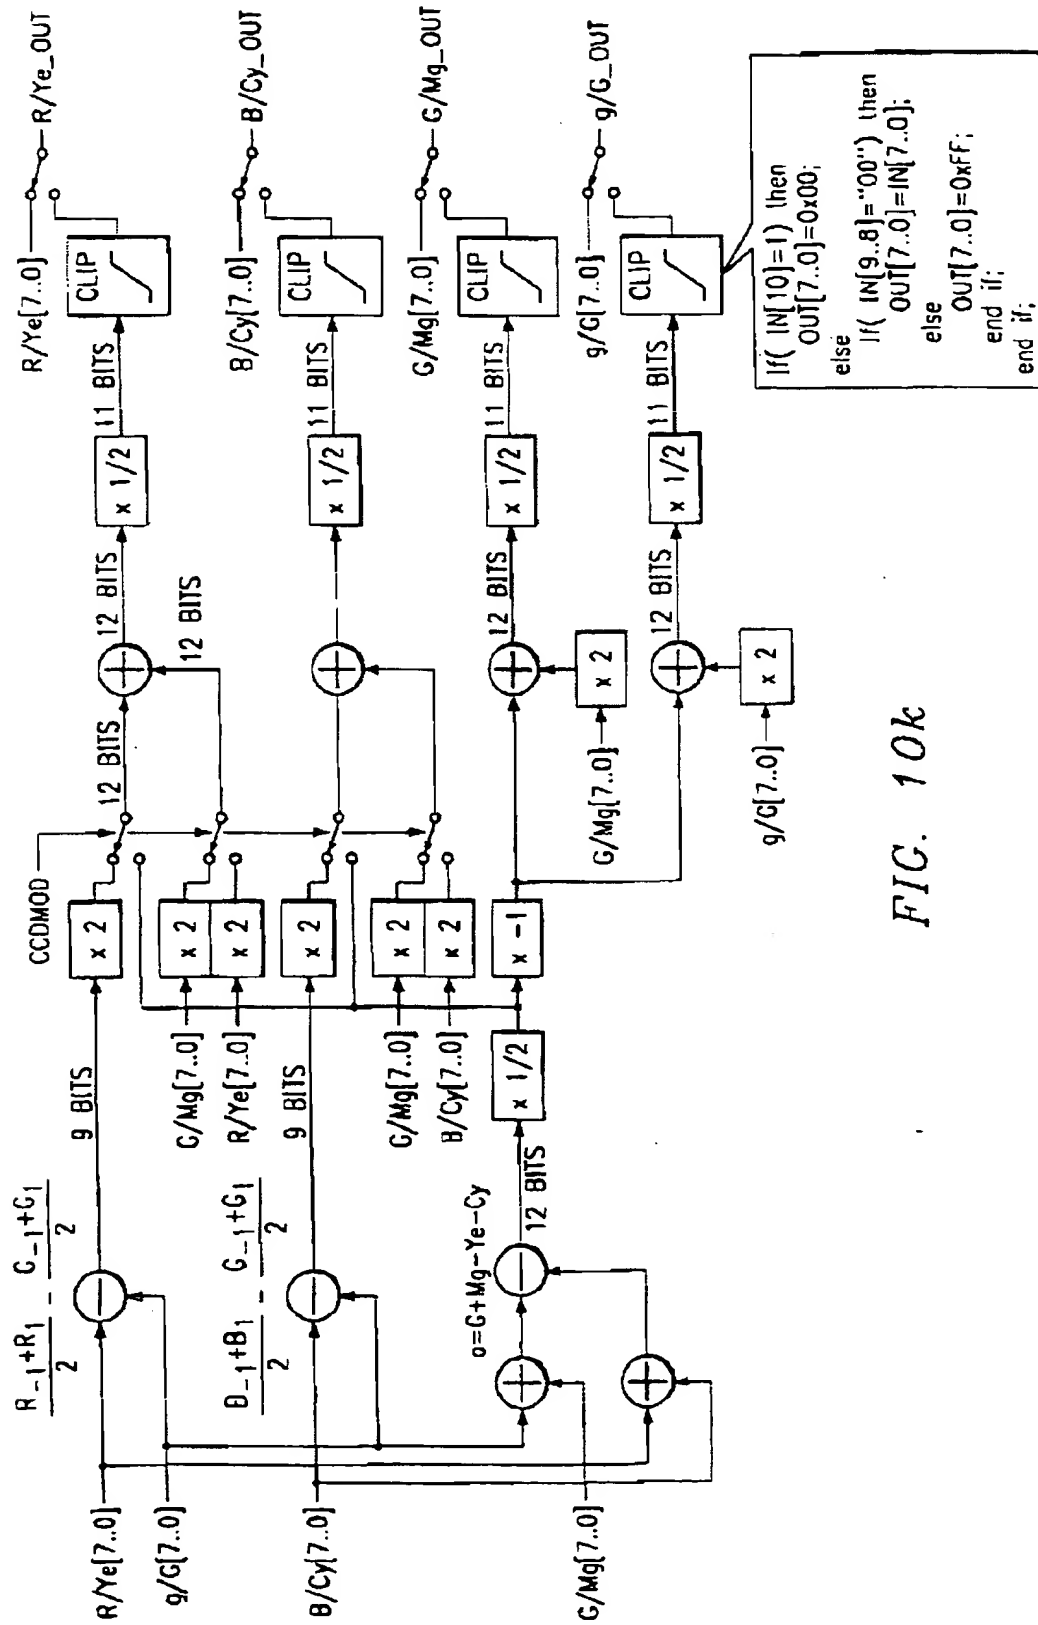


FIG. 10k

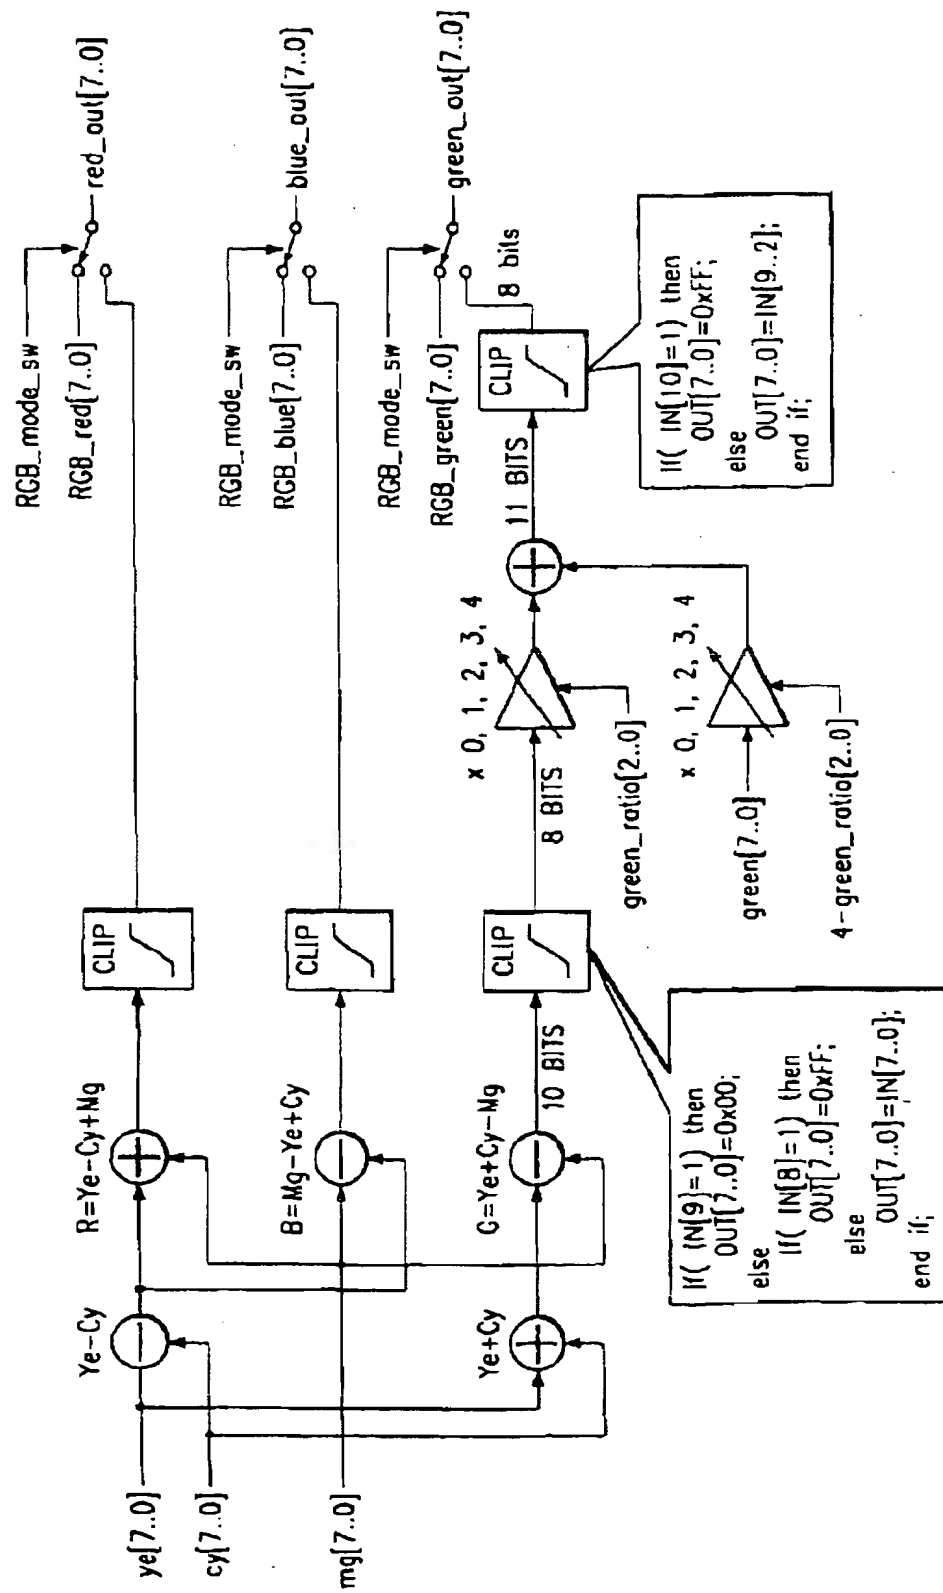


FIG. 10L

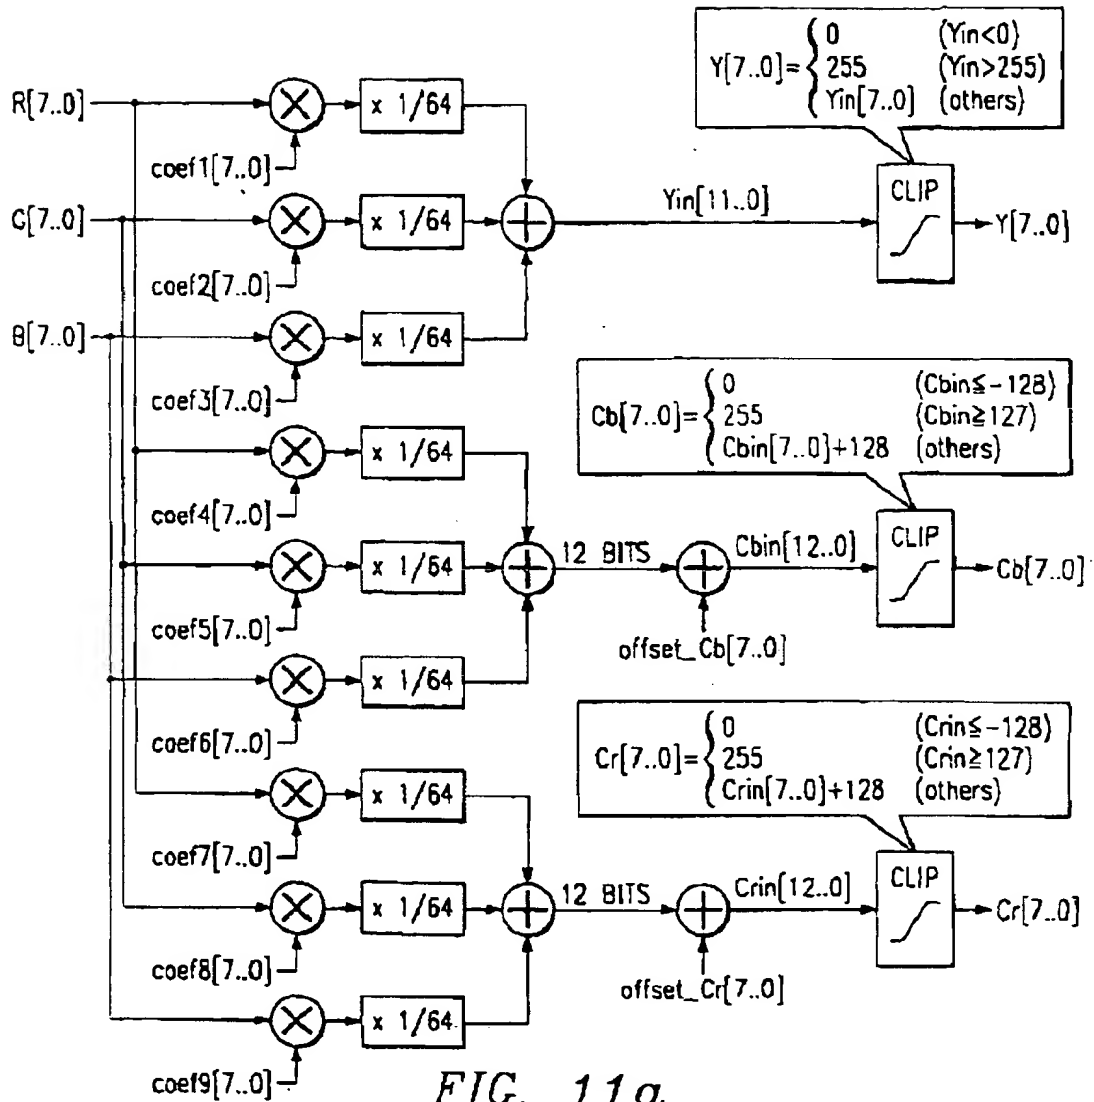


FIG. 11a


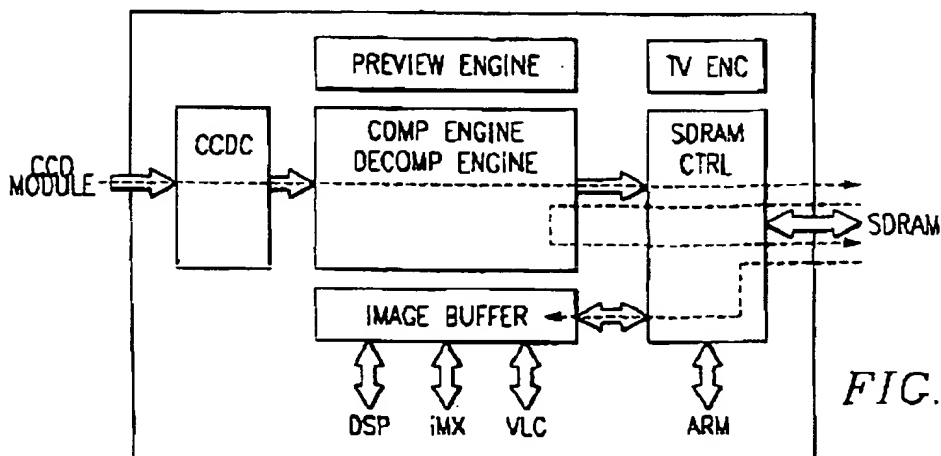
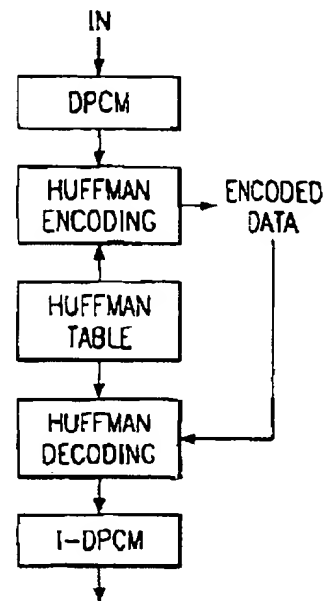
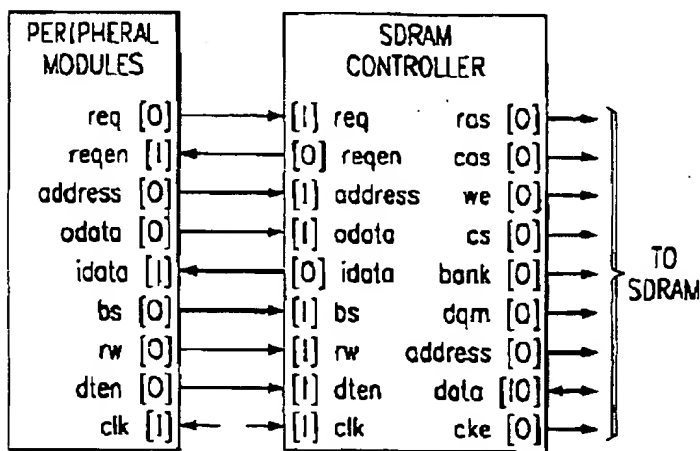
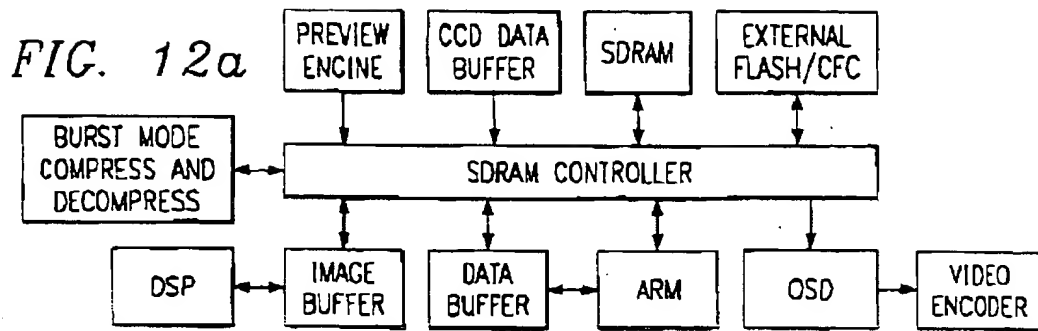




FIG. 11b








FIG. 16




FIG. 17

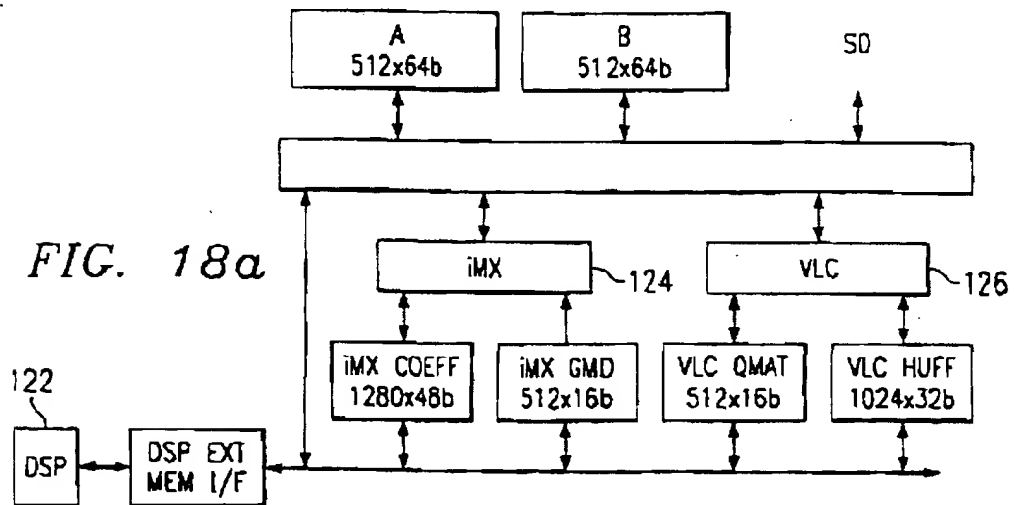


FIG. 18a




FIG. 18b




FIG. 19

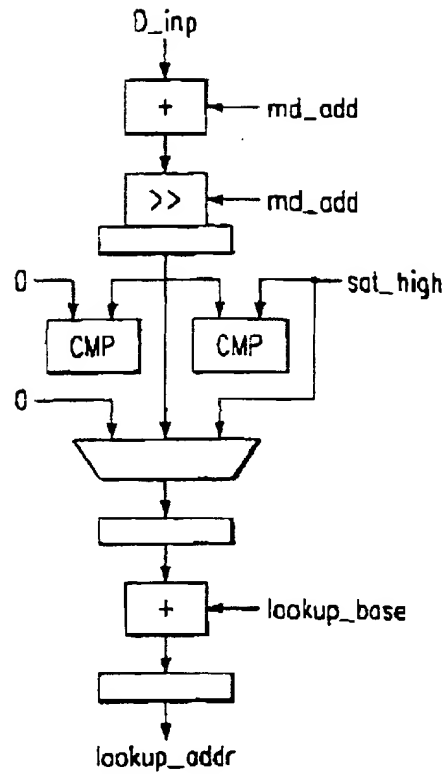


FIG. 21

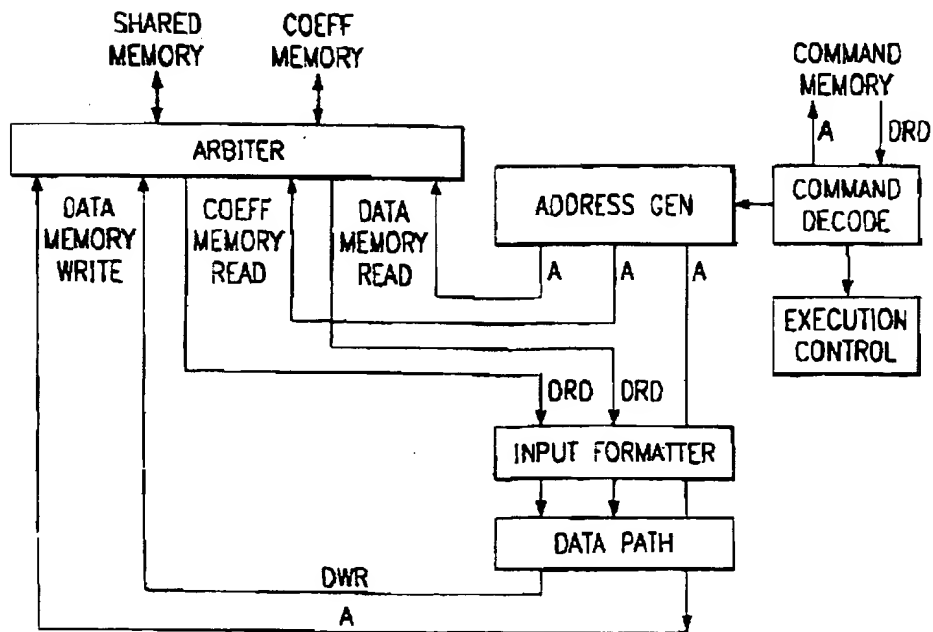


FIG. 20

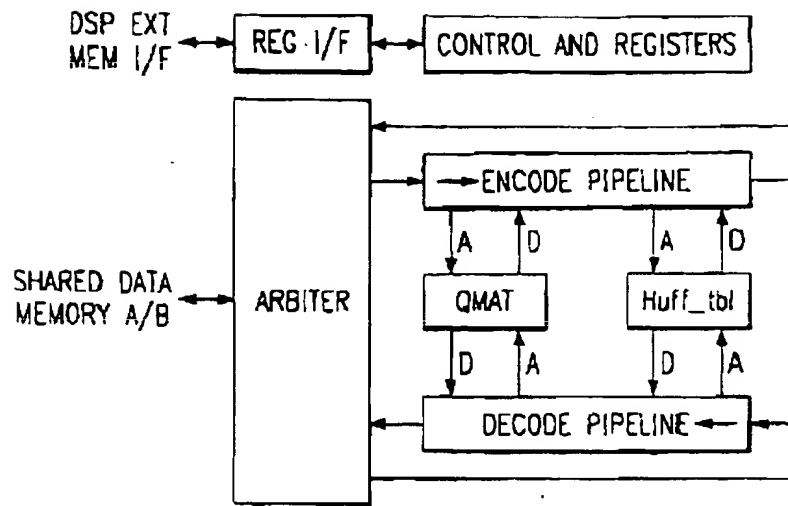


FIG. 22




FIG. 23a




FIG. 23b

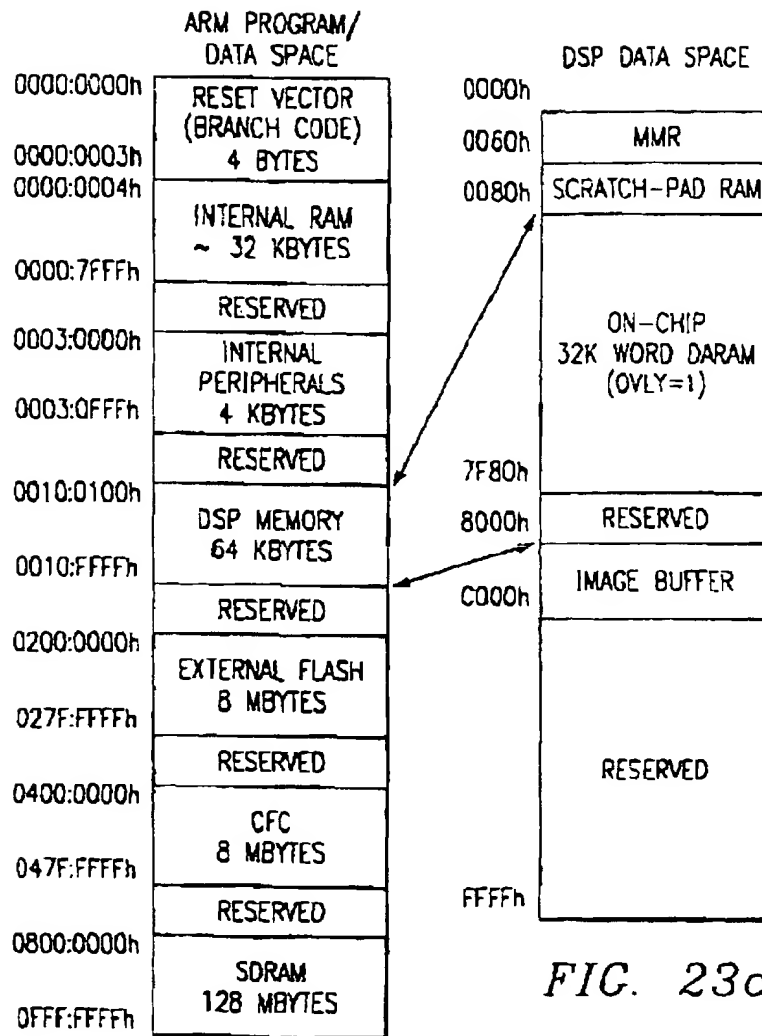


FIG. 23c

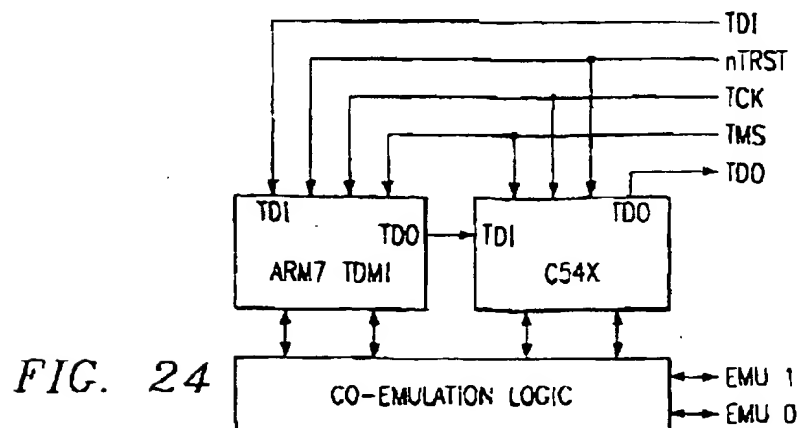


FIG. 24




FIG. 25




FIG. 26




FIG. 27




FIG. 28




FIG. 29




FIG. 30




FIG. 31




FIG. 32

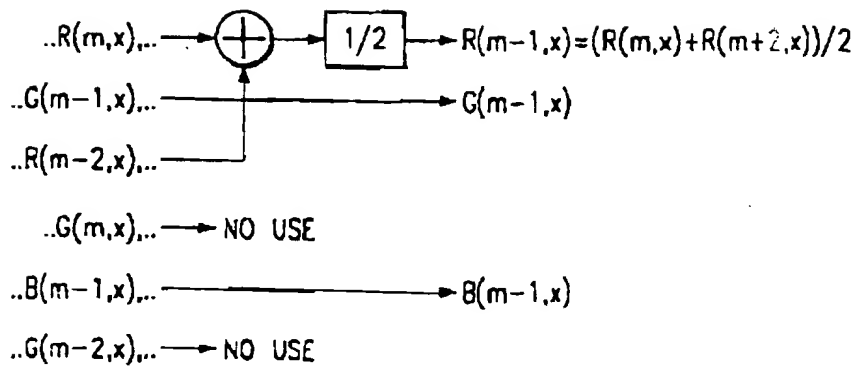


FIG. 33a

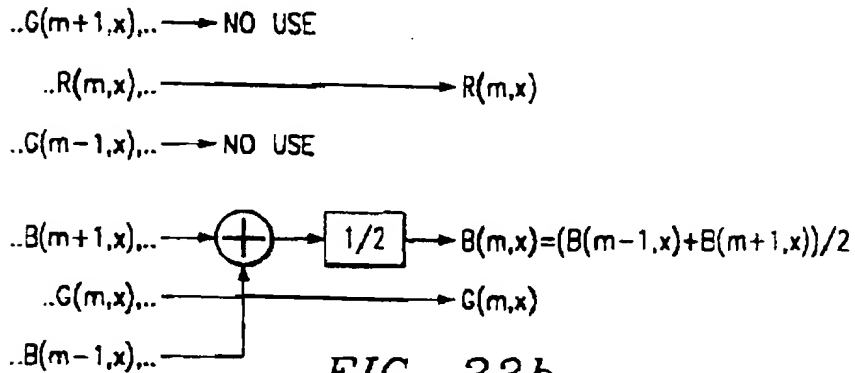


FIG. 33b

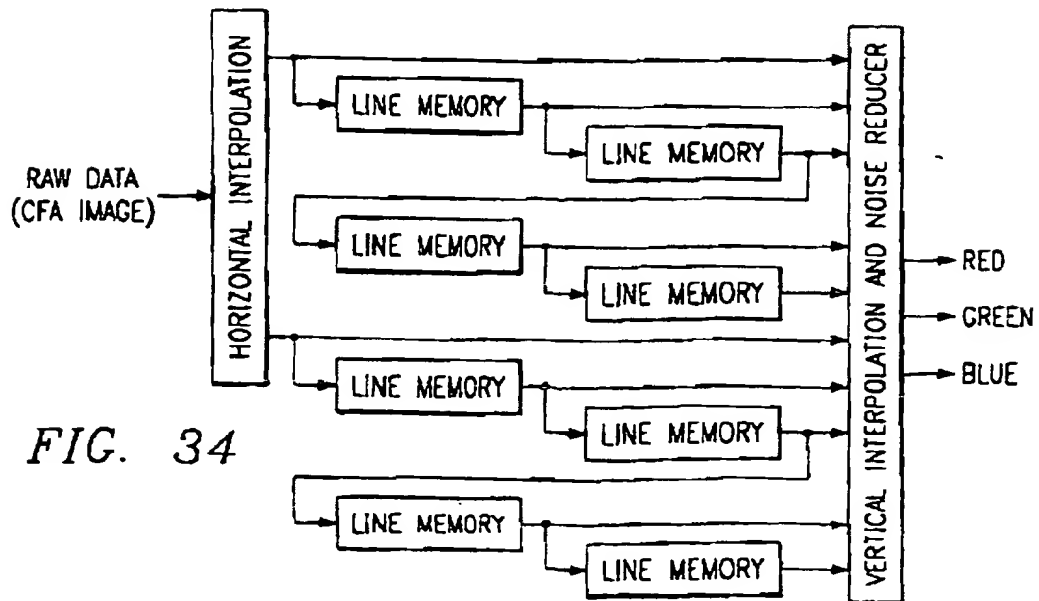



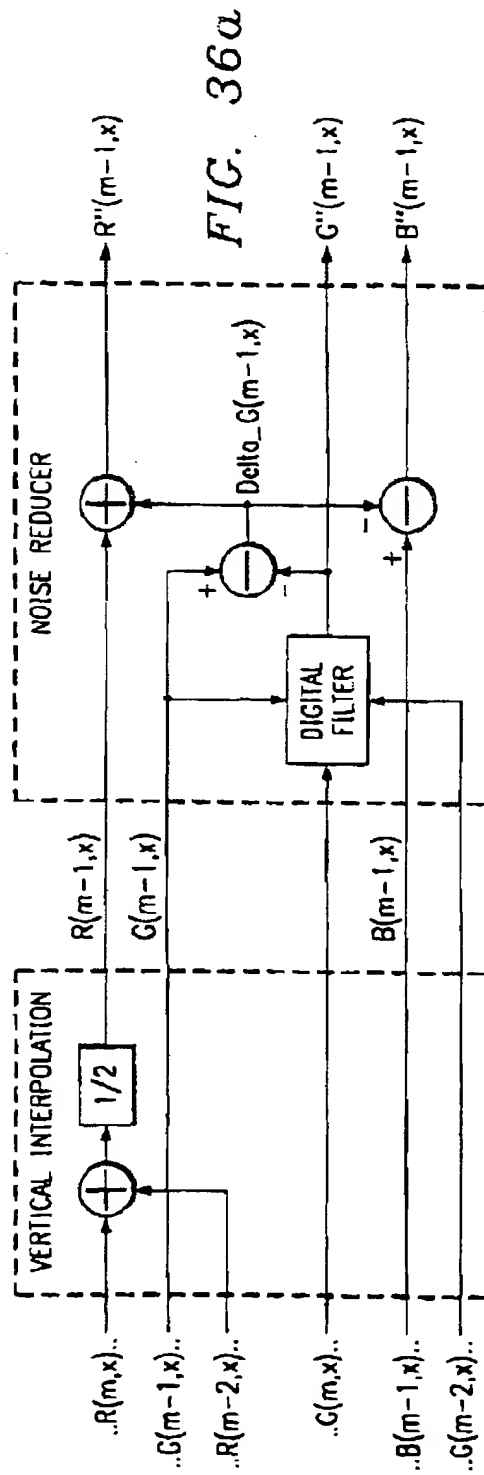





FIG. 34










FIG. 39a




FIG. 39b




FIG. 40




FIG. 41a




FIG. 41b




FIG. 42a




FIG. 42b






FIG. 42e






FIG. 43

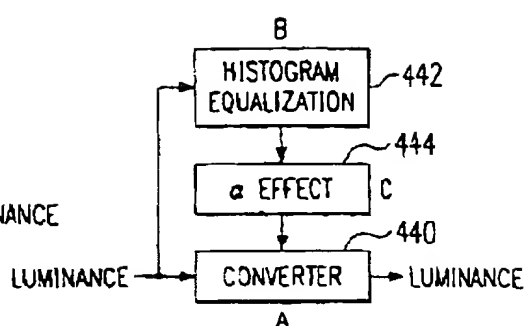


FIG. 44




FIG. 45




FIG. 46a




FIG. 46b

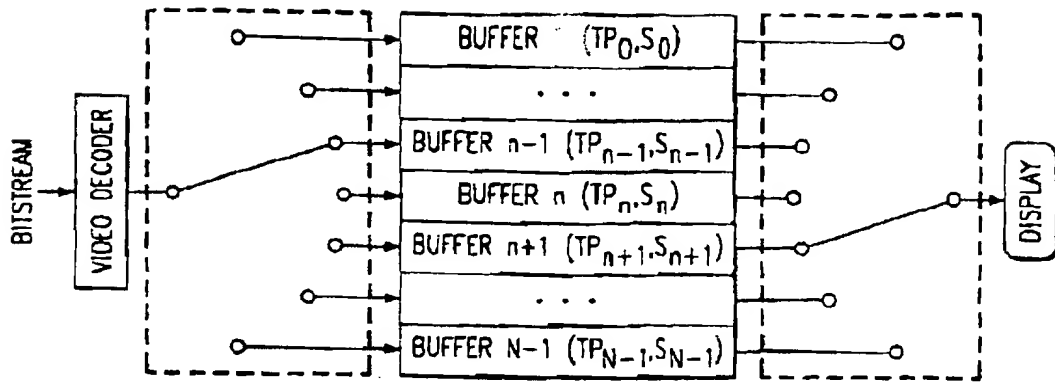


FIG. 47




FIG. 49




FIG. 50



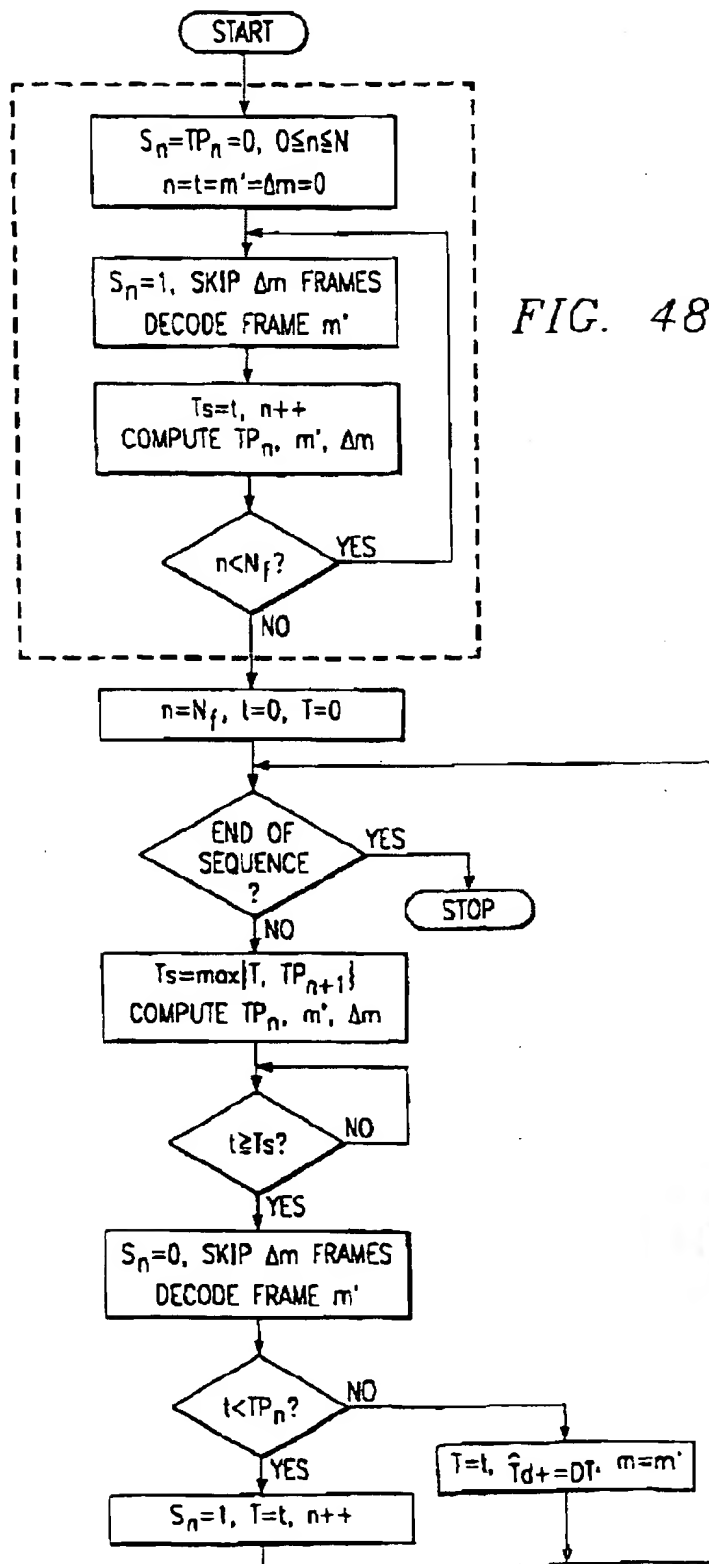


FIG. 51b






FIG. 52

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 111 904 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:

16.03.2005 Bulletin 2005/11

(51) Int Cl.7: H04N 1/40

(43) Date of publication A2:

27.06.2001 Bulletin 2001/26

(21) Application number: 00311430.3

(22) Date of filing: 20.12.2000

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR

Designated Extension States:

AL LT LV MK RO SI

(30) Priority: 20.12.1999 US 172780 P

29.06.2000 US 215000 P

29.06.2000 US 214951 P

04.08.2000 US 632543 P

14.01.2000 US 176272 P

(71) Applicant: Texas Instruments Incorporated

Dallas, Texas 75251 (US)

(72) Inventors:

- Hung, Ching-Yu
- Plano, Texas 75024 (US)

- Tamama, Hideo
- Tsukuba, Ibaraki (JP)

- Osamoto, Akira
- Inashiki, Ibaraki (JP)

- Koshiba, Osamu
- Tsukuba, Ibaraki (JP)

- Yamauchi, Satoru
- Tsuchiura, Ibaraki (JP)
- Zhou, Minhua
- Plano, Texas 75093 (US)
- Ilgner, Klaus
- 81549 München (DE)
- Talluri, Rajendra
- Plano, 75203 (US)
- Yoo, Youngjun
- Richardson, Texas 75080 (US)
- Liang, Jie
- Plano, Texas 75025 (US)
- Tsai, Mandy
- Taipei 111 (KR)
- Tsunoda, Kiyoshi
- Tokohama-city, Kanagawa (JP)
- Inamori, Shinri
- Yokohama-shi, Kanagawa (JP)

(74) Representative: Potter, Julian Mark et al

D Young & Co

120 Holborn

London EC1N 2DY (GB)

(54) Digital camera system and method

(57) Digital Still Camera (DSC) includes separate preview engine, burst mode compression/decompression engine, image pipeline, CCD plus CCD controller, and memory plus memory controller. ARM microprocessor and DSP share control. Color filter array interpolation by use of green high-frequency added to red and blue plus interpolation.




FIG. 1a

EP 1 111 904 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 00 31 1430

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	EP 0 675 648 A (EASTMAN KODAK CO) 4 October 1995 (1995-10-04) * abstract; figures 1,2 * -----	1	H04N1/40
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			H04N
<p>1 The present search report has been drawn up for all claims</p>			
Place of search		Date of completion of the search	Examiner
The Hague		5 October 2004	Isa, S
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			



European Patent
Office

Application Number
EP 00 31 1430

CLAIMS INCURRING FEES

The present European patent application comprised at the time of filing more than ten claims.

- ☐ Only part of the claims have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims and for those claims for which claims fees have been paid, namely claim(s):
- ☐ No claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims.

LACK OF UNITY OF INVENTION

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

see sheet B

- ☐ All further search fees have been paid within the fixed time limit. The present European search report has been drawn up for all claims.
- ☐ As all searchable claims could be searched without effort justifying an additional fee, the Search Division did not invite payment of any additional fee.
- ☐ Only part of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the inventions in respect of which search fees have been paid, namely claims:
- ☒ None of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the invention first mentioned in the claims, namely claims:

1



European Patent
Office

LACK OF UNITY OF INVENTION
SHEET B

Application Number
EP 00 31 1430

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

1. claim: 1

Integrated circuit for a digital still camera comprising
three programmable processors

2. claims: 2, and claims 12-14 as far as dependent on claim 2

A method for resizing a digital image

3. claims: 3, and claims 12-14 as far as dependent on claim 3

A method of colour-filtered array interpolation

4. claims: 4, and claims 12-14 as far as dependent on claim 4

A method of tone-scaling an image

5. claims: 5, and claims 12-14 as far as dependent on claim 5

A method of interpolation for a Bayer colour-filtered array

6. claims: 6, and claims 12-14 as far as dependent on claim 6

A method of interpolation for a Bayer colour-filtered array
image

7. claims: 7, and claims 12-14 as far as dependent on claim 7

A method of white balancing for a colour image sensor

8. claims: 8, and claims 12-14 as far as dependent on claim 8

A method of video decoding and display

9. claims: 9, and claims 12-14 as far as dependent on claim 9

A method of interpolation for a
complementary-colour-filtered array image

10. claims: 10, and claims 12-14 as far as dependent on claim 10

A method of interpolation for a colour-filtered array image



European Patent
Office

LACK OF UNITY OF INVENTION
SHEET B

Application Number
EP 00 31 1430

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

11. claims: 11, and claims 12-14 as far as dependent on claim 11

A method of playback buffering

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 00 31 1430

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

05-10-2004

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0675648 A	04-10-1995	US 5477264 A	19-12-1995
		DE 69522597 D1	18-10-2001
		DE 69522597 T2	11-07-2002
		EP 1152589 A2	07-11-2001
		EP 1152590 A2	07-11-2001
		EP 0675648 A2	04-10-1995
		JP 3597249 B2	02-12-2004
		JP 7274060 A	20-10-1995
		JP 2004048802 A	12-02-2004

EPO FORM P0429

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82



(12)

CORRECTED EUROPEAN PATENT APPLICATION

Note: Bibliography reflects the latest situation

(15) Correction information:

Corrected version no 1 (W1 A2)

IN/D code(s) 30

(51) Int Cl.⁷: **H04N 1/40**

(48) Corrigendum issued on:

16.08.2001 Bulletin 2001/33

(43) Date of publication:

27.06.2001 Bulletin 2001/26

(21) Application number: 00311430.3

(22) Date of filing: 20.12.2000

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR

Designated Extension States:

AL LT LV MK RO SI

(30) Priority: 20.12.1999 US 172780 P

29.06.2000 US 215000 P

29.06.2000 US 214951 P

04.08.2000 US 632543 P

14.01.2000 US 176272 P

(71) Applicant: **Texas Instruments Incorporated**
Dallas, Texas 75251 (US)

(72) Inventors:


- Hung, Ching-Yu
Plano, Texas 75024 (US)
- Tamama, Hideo
Tsukuba, Ibaraki (JP)
- Osamoto, Akira
Inashiki, Ibaraki (JP)
- Koshiba, Osamu
Tsukuba, Ibaraki (JP)

- Yamauchi, Satoru
Tsuchiura, Ibaraki (JP)
- Zhou, Minhua
Plano, Texas 75093 (US)
- Ilgner, Klaus
81549 München (DE)
- Talluri, Rajendra
Plano, 75203 (US)
- Yoo, Youngjun
Richardson, Texas 75080 (US)
- Liang, Jie
Plano, Texas 75025 (US)
- Tsai, Mandy
Taipei 111 (KR)
- Tsunoda, Kiyoshi
Tokohama-city, Kanagawa (JP)
- Inamori, Shinni
Yokohama-shi, Kanaqawa (JP)

(74) Representative: **Potter, Julian Mark et al**
D. Young & Co.,
21 New Fetter Lane
London EC4A 1DA (GB)

(54) **Digital camera system and method**

(57) Digital Still Camera (DSC) includes separate preview engine, burst mode compression/decompression engine, image pipeline, CCD plus CCD controller, and memory plus memory controller. ARM microprocessor and DSP share control. Color filter array interpolation by use of green high-frequency added to red and blue plus interpolation.



FIC. 1a